

Additive Spanners and (α, β) -Spanners

SURENDER BASWANA

Indian Institute of Technology Kanpur

TELIKEPALLI KAVITHA

Indian Institute of Science, Bangalore

KURT MEHLHORN

Max-Planck-Institut für Informatik

AND

SETH PETTIE

University of Michigan

5

Abstract. An (α, β) -spanner of an unweighted graph G is a subgraph H that distorts distances in G up to a multiplicative factor of α and an additive term β . It is well known that any graph contains a (multiplicative) $(2k - 1, 0)$ -spanner of size $O(n^{1+1/k})$ and an (additive) $(1, 2)$ -spanner of size $O(n^{3/2})$. However no other additive spanners are known to exist.

In this article we develop a couple of new techniques for constructing (α, β) -spanners. Our first result is an additive $(1, 6)$ -spanner of size $O(n^{4/3})$. The construction algorithm can be understood as an economical agent that assigns *costs* and *values* to paths in the graph, *purchasing* affordable paths and ignoring expensive ones, which are intuitively well approximated by paths already purchased. We show that this *path buying* algorithm can be parameterized in different ways to yield other sparseness-distortion tradeoffs. Our second result addresses the problem of which (α, β) -spanners can be computed efficiently, ideally in linear time. We show that, for any k , a $(k, k - 1)$ -spanner with size $O(kn^{1+1/k})$ can be found in linear time, and, further, that in a distributed network the algorithm terminates in a constant number of rounds. Previous spanner constructions with similar performance had roughly twice the multiplicative distortion.

This work was partially supported by the Future and Emerging Technologies program of the EU under contract number IST-1999-14186 (ALCOM-FT).

Authors' addresses: S. Baswana, Department of Computer Science and Engineering, IIT Kanpur, Kanpur 208016, India; e-mail: sbaswana@cse.iitk.ac.in; T. Kavitha, Department of Computer Science and Automation, Indian Institute of Science, Bangalore 560012, India; e-mail: kavitha@cse.iisc.ernet.in; K. Mehlhorn, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 61, 66123 Saarbrücken, Germany; e-mail: mehlhorn@mpi-inf.mpg.de; S. Pettie, Department of Electrical Engineering and Computer Science, University of Michigan, 2260 Hayward Street, Ann Arbor, MI 48109; e-mail: pettie@umich.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1549-6325/2010/11-ART5 \$10.00

DOI 10.1145/1868237.1868242 <http://doi.acm.org/10.1145/1868237.1868242>

Categories and Subject Descriptors: G.2.2 [Discrete Mathematics]: Graph Theory—*Graph algorithms*

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Spanner, metric embedding

ACM Reference Format:

Baswana, S., Kavitha, T., Mehlhorn, K., and Pettie, S. 2010. Additive spanners and (α, β) -spanners. ACM Trans. Algor. 7, 1, Article 5 (November 2010), 26 pages.
DOI = 10.1145/1868237.1868242 <http://doi.acm.org/10.1145/1868237.1868242>

1. Introduction

An (α, β) -spanner of an undirected graph G is a *subgraph* H such that for all vertices u, v :

$$\delta_H(u, v) \leq \alpha \cdot \delta_G(u, v) + \beta,$$

where δ_G is the distance in graph G . In other words, an (α, β) -spanner guarantees that for pairs of vertices far apart in G , their distance in the spanner is stretched by roughly an α factor, which would ideally be close to 1. We call a $(1, \beta)$ -spanner an *additive β -spanner*. If $\beta = 0$, this definition reverts to the usual definition of a *multiplicative α -spanner* [Peleg and Schaffer 1989; Althöfer et al. 1993].

Spanners (and related structures) are useful in many contexts. They are the basis of space-efficient routing tables that guarantee nearly shortest routes [Abraham et al. 2006b; Thorup and Zwick 2001; Roditty et al. 2002; Cowen 2001; Cowen and Wagner 2004; Peleg and Upfal 1989], schemes for simulating synchronized protocols in unsynchronized networks [Peleg and Ullman 1989], and parallel and distributed algorithms for computing approximate shortest paths [Cohen 1998, 2000; Elkin 2005]. A recent application of spanners is the construction of labeling schemes and distance oracles [Thorup and Zwick 2005; Baswana and Sen 2006; Roditty et al. 2005; Baswana and Kavitha 2006; Baswana et al. 2008], which are data structures that can report approximately accurate distances in constant time. In all of these applications, the quality of the solution ultimately depends on an efficient algorithm for computing a *low-distortion* sparse spanner. The main open problem in this area is to understand the inherent tradeoffs between these three measures of efficiency: distortion (α and β), sparseness, and construction time. Even ignoring construction time, there are only a handful of cases where the distortion-sparseness tradeoff is fully understood.

1.1. MULTIPLICATIVE SPANNERS. The early work on spanners established the basic tradeoff between sparseness and *multiplicative* distortion. If the spanner size is fixed at $O(n^{1+1/k})$, the multiplicative distortion can be no better than $\Theta(k)$ [Peleg and Schaffer 1989]. We let n and m be the number of vertices and edges in the input graph. Althöfer et al. [1993] proposed a greedy algorithm for producing a $(2k - 1)$ -spanner whose size is at most $m_{2k+1}(n)$, where $m_g(n)$ is the maximum number of edges in a graph with girth at least g .¹ Moreover, they observed that $m_{2k+1}(n)$ is *precisely* the best possible bound for a $(2k - 1)$ -spanner. If one removes

¹ Girth is the length of the shortest cycle. Note that, since every graph has a bipartite subgraph with at least half the edges, $\frac{1}{2}m_{2k+1}(n) \leq m_{2k+2}(n) \leq m_{2k+1}(n)$.

any edge from a graph with girth $2k + 1$, the distance between its endpoints jumps from 1 to at least $2k$. Thus, the only $(2k - 1)$ -spanner of such a graph is the graph itself. A trivial upper bound on $m_{2k+1}(n)$ and $m_{2k+2}(n)$ is $O(n^{1+1/k})$. It has been conjectured, by Erdős[1963] and others that this bound is asymptotically tight, though the conjecture has only been proved for $k = 1, 2, 3$, and 5 ; weaker lower bounds are known for all other k ; see Wenger [1991] and Thorup and Zwick [2005]. In other words, finding the exact tradeoff between sparseness and *multiplicative* distortion is at least as hard as proving or disproving the girth conjecture.

The best-known implementations of the Althöfer et al. algorithm run in time $O(\min\{kn^{2+1/k}, mn^{1+1/k}\})$ [Althöfer et al. 1993; Roditty and Zwick 2004], though there are several more efficient $(2k - 1)$ -spanner constructions. Halperin and Zwick [1996] (see Peleg [2000]) computed an $O(n^{1+1/k})$ -size $(2k - 1)$ -spanner in linear time. However, unlike the algorithm of Althöfer et al. [1993], the Halperin-Zwick algorithm only works on unweighted graphs. For weighted graphs Baswana and Sen [2007] gave a randomized construction of such a spanner with size $O(kn^{1+1/k})$. The Baswana-Sen algorithm has since been derandomized by Roditty et al. [2005].

1.2. BEYOND PURELY MULTIPLICATIVE DISTORTION. The *girth* bound exactly characterizes the optimal tradeoff between sparseness and multiplicative distortion but arguments based on girth only apply to *adjacent* vertices. In unweighted graphs, the girth argument could just as easily be interpreted as bounding the additive distortion, or some combination of additive and multiplicative distortion. In particular, if the girth conjecture is true, we can only say that an (α, β) -spanner of size $O(n^{1+1/k})$ has $\alpha + \beta \geq 2k - 1$. It is conceivable that there exist additive $(2k - 2)$ -spanners with size $O(n^{1+1/k})$, for any k . Before our work, however, only one such additive spanner was known. Aingworth et al. [1999] (with followup work in Dor et al. [2000]; Elkin and Peleg [2004]; Thorup and Zwick [2006], and Roditty et al. [2005]) showed that there exist $O(n^{3/2})$ -size additive 2-spanners. On the lower-bound side, Woodruff [2006] recently proved that any spanner with size $O(k^{-1}n^{1+1/k})$ cannot do better than an additive distortion of $2k - 2$, *independent* of whether the girth conjecture is true or not.

The current research trend is to optimize distortion as a function of the distance being approximated, rather than fixate on adjacent vertices and the girth conjecture. Elkin and Peleg [2001, 2004] were the first to take this approach in the context of general graphs. They showed the existence of $(k - 1, O(k))$ -spanners with size $O(kn^{1+1/k})$ and $(1 + \epsilon, \beta_{k,\epsilon})$ -spanners with size $O(\beta_{k,\epsilon}n^{1+1/k})$, where $\beta_{k,\epsilon}$ is roughly $(\epsilon^{-1} \log k)^{\log k}$. Elkin and Peleg [2004] and Elkin [2005] gave algorithms for constructing $(1 + \epsilon, \beta_{k,\epsilon,t})$ -spanners in time, respectively, $O(n^{2+1/t})$ and $O(mn^{\frac{1}{2k} + \frac{1}{t}})$, where $\beta_{k,\epsilon,t}$ can be exponentially larger than $\beta_{k,\epsilon}$. Following Thorup and Zwick [2006], Pettie [2007] gave $(1 + \epsilon, \beta)$ -spanners where ϵ is *not* a parameter of the construction but can be chosen as a function of the distance d being approximated. In particular, for any $o \geq 0$ there is a spanner with size $O(n^{1 + \frac{(3/4)^o}{7 - 2(3/4)^o}})$ such that vertices at distance d appear at distance $d + O(d^{1-1/(o+2)} + o^o)$ in the spanner. Pettie [2007] also showed that every graph has a $(1 + \epsilon, \beta)$ -spanner with size $O(n \log \log(\epsilon^{-1} \log \log n))$, where $\beta = O(\epsilon \log \log n)^{\log \log n}$. The construction algorithms of Elkin and Peleg [2001, 2004], Elkin [2005], and Thorup and Zwick [2006] are faster than that of Pettie [2007] though the resulting spanners have higher distortion. See Figure 1 for a tabular summary of existing spanners constructions.

WEIGHTED GRAPHS, MULTIPLICATIVE SPANNERS

α	Size	Time	Notes
$2k - 1$	$\left\{ \begin{array}{l} \frac{1}{2}n^{1+1/k} \\ O(kn^{1+1/k}) \end{array} \right.$	$\left\{ \begin{array}{l} O(mn^{1+1/k}) \\ O(kn^{2+1/k}) \end{array} \right.$	Althöfer et al. [1993] Roditty et al. [2004]
		$\left\{ \begin{array}{l} O(km) \text{ (rand.)} \\ O(km) \text{ (det.)} \end{array} \right.$	Baswana and Sen [2007] Roditty et al. [2005]

UNWEIGHTED GRAPHS, (α, β) -SPANNERS

(α, β)	Size	Time	Notes
$(2k - 1, 0)$	$n^{1+1/k}$	$O(m)$	see Peleg [2000]
$(k - 1, O(k))$	$O(kn^{1+1/k})$	$O(mn^{1-1/k})$	Elkin and Peleg [2001]
$(k, k - 1)$	$O(kn^{1+1/k})$	$O(km)$	new
$(1 + \epsilon, 4)$	$\left\{ \begin{array}{l} O(\epsilon^{-1}n^{4/3}) \\ O(\beta n^{1+1/k}) \end{array} \right.$	$O(mn^{2/3})$	Elkin and Peleg [2004]
		$\left\{ \begin{array}{l} O(n^{2+1/t}) \\ O(mn^{\frac{1}{2k} + \frac{1}{t}}) \end{array} \right.$	Elkin and Peleg [2004] Elkin [2005]
$(1 + \epsilon, \beta)$	$\left\{ \begin{array}{l} O(kn^{1+1/k}) \\ O(on^{1 + \frac{(3/4)\sigma}{7-2(3/4)\sigma}}) \\ O(n \log \log \epsilon^{-1}) \end{array} \right.$	$O(kmn^{1/k})$	Thorup and Zwick [2006]
		poly(n)	Pettie [2007]
		poly(n)	Pettie [2007]
$\left. \begin{array}{l} (1, n^{\frac{9}{16} - \frac{7\delta}{8}}) \\ (1, n^{1-3\delta}) \end{array} \right\}$	$O(n^{1+\delta})$	poly(n)	Pettie [2007] new
$(1, 2)$	$O(n^{3/2})$	$\left\{ \begin{array}{l} O(m\sqrt{n}) \\ \tilde{O}(n^2) \end{array} \right.$	Aingworth et al. [1999] Dor et al. [2000], Roditty et al. [2005]
$(1, 6)$	$O(n^{4/3})$	$O(mn^{2/3})$	new

FIG. 1. The state-of-the-art in (α, β) -spanners. The size bound of the spanners of Althöfer et al. [1993] and Roditty and Zwick [2004] follow from results by Alon et al. [2002]. The $(1 + \epsilon, \beta)$ -spanners vary widely in β 's dependence on ϵ and whether the distortion holds for a specific ϵ or *all* ϵ simultaneously. The sparseness-distortion guarantees of Pettie [2007] dominate those of Elkin and Peleg [2004], Elkin [2005], and Thorup and Zwick [2006], though the construction algorithms of Elkin and Peleg [2004], Elkin [2005], and Thorup and Zwick [2006] are faster. In Pettie [2007] and Thorup and Zwick [2006], ϵ can be chosen as a function of the distance d being approximated. If ϵ is chosen optimally, the distance between two vertices in the spanner of Pettie [2007] is $d + O(d^{1-1/(o+2)} + o^o)$, where $0 \leq o < \log_{4/3} \log n$ controls the size of the spanner. This is always better than the bound of Thorup and Zwick [2006]: $d + O(d^{1-1/(k-1)} + 2^k)$, where $2 \leq k < \log n$. The β 's from Elkin and Peleg [2004], Elkin [2005], and Pettie [2007] are, respectively, $\beta \approx \max\{\epsilon^{-1} \log k, t\}^{\log k}$, $\beta \approx (\epsilon^{-1} t^2 k \log k)^{t \log k}$, and $\beta \approx (\epsilon^{-1} \log \log n)^{\log \log n}$.

1.3. OUR RESULTS. Our first result is that every graph contains an additive 6-spanner with size $O(n^{4/3})$ and that such a spanner can be computed efficiently. This result is a far cry from a full spectrum of tradeoffs between sparseness and additive distortion. However, our approach is completely new and is generic enough to be applied in other ways. We view a spanner construction as an economic agent that assigns a *cost* and *value* to paths in the graph. Affordable paths are *purchased* (included in the spanner) and expensive ones ignored. Different cost and value combinations lead to spanners with different properties. Besides constructing a 6-spanner, our path buying algorithm can be parameterized to find an additive 2-spanner of size $O(n^{3/2})$, matching Aingworth et al. [1999], Dor et al. [2000],

Elkin and Peleg [2004], Thorup and Zwick [2006], and Roditty et al. [2005], and an additive $(n^{1-3\delta})$ -spanner with size $O(n^{1+\delta})$, for any constant $\delta \in (0, 1/3)$. The latter result improves on the additive $\tilde{O}(n^{9/16-7\delta/8})$ -spanners of Pettie [2007] for some δ . Pettie [2007], following the initial publication of our work [Baswana et al. 2005], used a hierarchical version of the path buying algorithm to construct sparser $(1 + \epsilon, \beta)$ -spanners.

Our second result addresses those sparseness-distortion tradeoffs that can be computed by an *efficient* algorithm. We show that a $(k, k - 1)$ -spanner of size $O(kn^{1+1/k})$ can be constructed in $O(km)$ time. Since the decisions made by the algorithm are very local, it can easily be implemented in modern models of computation. For instance, in the cache-oblivious model [Frigo et al. 1999], the PRAM model [Karp and Ramachandran 1990], or in a synchronized distributed network [Peleg 2000], our algorithm is close to optimal under the relevant measures. Previous spanners with equal or better distortion [Elkin and Peleg 2004; Elkin 2005; Thorup and Zwick 2006] have construction times of the form $O(mn^{\Omega(1)})$ and the result that is most comparable to ours, the $(k - 1, O(k))$ -spanner of Elkin and Peleg [2001], requires time $O(mn^{1-1/k})$ to compute. If we restrict our attention to near-linear time constructions, all the existing spanners with size $O(kn^{1+1/k})$ [Halperin and Zwick 1996; Baswana and Sen 2007; Roditty et al. 2005] had multiplicative distortion $2k - 1$. Whereas our $(k, k - 1)$ -spanners can be computed in $O(k)$ rounds in a distributed network, all previous constructions with equal or better distortion required $\Omega(n^{\Omega(1)})$ rounds [Elkin and Zhang 2006].

1.4. RELATED WORK. Spanners are part of a large body of work on *metric embeddings*, where one wants a mapping $\phi : S \rightarrow T$ from a given (finite) source metric² (S, δ_S) to a target metric (T, δ_T) that does not distort interpoint distances by too much. (Distortion here is usually defined as purely multiplicative distortion.) In our case (S, δ_S) is some unweighted graph metric and ϕ is the identity function; the problem is to find a metric (T, δ_T) corresponding to a sparse subgraph. We only review metric embedding results where the target metric is some kind of graph. See Indyk [2001] and Indyk and Matousek [2004] for surveys of metric embedding theory.

Roditty et al. [2002] constructed multiplicative spanners for directed graphs under the *roundtrip* metric, where the distance between two vertices is the length of the shortest cycle (not necessarily simple) that contains both. The sparseness-distortion tradeoffs in Roditty et al. [2002] are slightly worse than those obtained for undirected graphs. Bollobás et al. [2003] and Coppersmith and Elkin [2005] studied spanners that preserve, without distortion, the distance between some pairs of vertices. Bollobás et al. [2003] showed that there exist $O(n^{1+\delta})$ -size spanners that preserve distances greater than $n^{1-\delta}$, and that this tradeoff is optimal. Coppersmith and Elkin [2005] showed that for any set P of pairs of vertices, there is a distance preserver for P with size $O(n + |P| \sqrt{n})$ edges, and that this bound is optimal in some circumstances. In particular, if $|P| = \omega(\sqrt{n})$ then a size of $O(n)$ cannot be guaranteed. Dor et al. [2000] considered *emulators*, which may contain both graph edges and additional weighted edges, and observed that there are additive 4-emulators

²Recall that (S, δ_S) is a metric if for $u, v, w \in S$, $\delta_S(u, u) = 0$, $\delta_S(u, v) = \delta_S(v, u) \geq 0$, and $\delta_S(u, v) \leq \delta_S(u, w) + \delta_S(w, v)$.

with $O(n^{4/3})$ edges. Thorup and Zwick [2006] generalized this construction to emulators with $O(kn^{1+\frac{1}{2k-1}})$ edges and additive distortion $O(d^{1-\frac{1}{k-1}} + k^k)$, where d is the distance being approximated. One well-known application of (multiplicative) spanners is in the construction of approximate distance oracles; see Thorup and Zwick [2005], Baswana and Sen [2006], Roditty et al. [2005], Baswana and Kavitha [2006], Mendel and Naor [2007], and Baswana et al. [2008].

The sparsest spanner is a tree, but it is impossible to guarantee that a tree spanner has any nontrivial *worst-case* distortion.³ A number of weaker notions of distortion have been defined to deal with tree spanners. A *probabilistic embedding* with distortion t is a distribution over tree metrics such that $\mathbb{E}_T[\delta_T(\phi(u), \phi(v))/\delta_S(u, v)] \leq t$, where it is assumed that $\delta_T(\phi(u), \phi(v)) \geq \delta_S(u, v)$. Elkin et al. [2005] showed that for any graph there is a probabilistic embedding into its spanning trees with distortion $O(\log^2 n \log \log n)$. Fakcharoenphol et al. [2004] proved that *any* metric can be probabilistically embedded in a tree metric (not necessarily a spanning tree) with distortion $O(\log n)$. These probabilistic embeddings have surprisingly diverse applications; see Alon et al. [1995], Bartal [1996], Fakcharoenphol et al. [2004], Spielman and Teng [2004], and Xiao et al. [2007], and the references therein. The distinction between embedding into a spanning subtree versus any tree metric was explored by Bădoiu et al. [2007], who showed that the distortion of the best spanning subtree is between $\Omega(\log n / \log \log n)$ and $O(\log n)$ times the distortion of the best tree metric. They also give an algorithm to approximate the best embedding from a given metric into a tree metric.⁴

A number of recent articles have looked at spanners with ϵ -slack, meaning the stated distortion (a function of ϵ) may fail to hold for an ϵ fraction of the vertex pairs. Such a spanner is *gracefully degrading* if it has ϵ -slack for all ϵ . Chan et al. [2006] gave a linear-size, gracefully degrading spanner with multiplicative distortion $O(\log \epsilon^{-1})$. See Abraham et al. [2005, 2006a, 2007] for standard and probabilistic embeddings into tree metrics with ϵ -slack.

For geometric graphs, where the vertices are points in \mathbb{R}^d , it is known that, for any constants d, ϵ , there are efficiently constructible linear size $(1 + \epsilon)$ -spanners [Gudmundsson et al. 2002; Narasimhan and Smid 2007]. Geometric graphs fall into a larger class of metrics with constant *doubling* dimension.⁵ It was recently shown that even these metrics have $(1 + \epsilon)$ -spanners with size $O(n)$, for constant ϵ and dimension d [Chan et al. 2005; Chan and Gupta 2006; Har-Peled and Mendel 2006].

1.5. ORGANIZATION. In Section 2 we introduce the path-buying technique and an algorithm for finding additive 6-spanners. In Section 2.1 we show how the path-buying algorithm can be parameterized to compute other additive spanners. Section 3 contains our linear time algorithm for constructing $(k, k - 1)$ -spanners and in Section 3.1 we show how it can easily be adapted to other models of computation. In Section 4 we conclude the article with some open problems.

³For example, consider a cycle of n vertices.

⁴Notice the difference between absolute versus relative distortion. Most results give an absolute guarantee on the distortion (perhaps in expectation) whereas Bădoiu et al. [2007] compared the distortion of their embedding against the optimal one. See Emek and Peleg [2004] and Elkin and Peleg [2005] for other (in)approximability results for different spanner problems.

⁵A metric has doubling dimension d if the ball of radius $2r$ centered at any point can be covered by 2^d balls of radius r .

2. Additive Spanners

Our construction of additive 6-spanners works in two phases, the first of which involves standard clustering techniques. For the sake of completeness we will describe a simple randomized clustering algorithm developed by Baswana and Sen [2007]. A more general version of this algorithm appears in Section 3, as well as a deterministic alternative.

In phase one, we choose a set of $n^{2/3}$ vertices at random to become *cluster centers*.⁶ Every vertex will join the cluster of an arbitrary adjacent center and, if there are no adjacent centers, the vertex will be left *unclustered*. Let $\mathcal{C} = \{C_1, C_2, \dots, C_{n^{2/3}}\}$ be the set of clusters; that is, C_i consists of the i th center and all adjacent vertices that joined it. The edge set H_0 (which is a subset of our spanner) consists of a radius-1 spanning tree of each cluster and *all* edges that are incident to at least one unclustered vertex. Consider the number of edges contributed to H_0 by some vertex v with degree $\deg(v)$. If v is clustered, but not a center, it contributes one edge to the spanning tree of its cluster. This happens with probability *at least* $1 - (1 - n^{-1/3})^{\deg(v)+1}$. With probability at most $(1 - n^{-1/3})^{\deg(v)+1}$, v is left unclustered and contributes $\deg(v)$ edges. Thus the expected contribution from v is at most $1 + (\deg(v) - 1)(1 - n^{-1/3})^{\deg(v)+1}$, which is always less than $n^{1/3}$. By linearity of expectation, $\mathbb{E}[|H_0|] < n^{4/3}$. See Section 3 for a more detailed analysis.

Since H_0 contains all edges incident to unclustered vertices, we can focus our attention on shortest paths whose endpoints are both clustered. The objective of phase two is to show that on any shortest path $P = \langle u, \dots, u' \rangle$, where both u and u' are clustered, there exists a short path Q in the spanner from u to u' that passes through some $C^* \in \mathcal{C}(P)$. Here $\mathcal{C}(X)$ represents the set of clusters intersecting X or the unique cluster containing X if X is a single vertex. We guarantee, in particular, that the portions of Q from $\mathcal{C}(u) \rightsquigarrow C^*$ and $C^* \rightsquigarrow \mathcal{C}(u')$ are no longer than their counterparts in P . Property 2.1 formalizes this idea, and Lemma 2.2 states that any subgraph with this property is an additive 6-spanner.

PROPERTY 2.1. *A subgraph $H \supseteq H_0$ is happy if for any two clustered vertices u, u' , there exists a shortest path $P = \langle u, \dots, u' \rangle$ in G and a $C^* \in \mathcal{C}(P)$ such that*

$$\delta_H(\mathcal{C}(v), C^*) \leq \delta_P(\mathcal{C}(v), C^*) \quad \text{for both } v \in \{u, u'\}.$$

LEMMA 2.2. *Any happy subgraph of G is also an additive 6-spanner of G .*

PROOF. Let H be the happy subgraph, and u, u', P , and $C^* \in \mathcal{C}(P)$ be as in the statement of Property 2.1; see Figure 2. We can bound the distance from u to u' in H as

$$\begin{aligned} \delta_H(u, u') &\leq \text{diam}_H(\mathcal{C}(u)) + \delta_H(\mathcal{C}(u), C^*) + \text{diam}_H(C^*) \\ &\quad + \delta_H(C^*, \mathcal{C}(u')) + \text{diam}_H(\mathcal{C}(u')) \\ &\leq \delta_P(\mathcal{C}(u), C^*) + \delta_P(C^*, \mathcal{C}(u')) + 6 \\ &\leq \delta_G(u, u') + 6, \end{aligned}$$

where $\text{diam}_H(Z)$ represents the maximum distance between vertices in Z in the subgraph H . The second inequality follows directly from Property 2.1. \square

⁶ Sampling each vertex independently with probability $n^{-1/3}$ achieves the same goal.




FIG. 2. The clusters $\mathcal{C}(u)$, \mathcal{C}^* , and $\mathcal{C}(u')$ indicated by ovals. The shortest intercluster paths in H are indicated by dashed curves.

In phase two, we find a subgraph $H_0 \cup P_1 \cup P_2 \cup \dots \cup P_p$ where P_1, \dots, P_p are paths purchased by the *path buying* algorithm in Figure 3. The algorithm is parameterized by *cost* and *value* functions. It evaluates some shortest path between each pair of vertices and purchases the path if twice its value exceeds its cost. If the path is bought, this will influence the cost and value of other paths.

The following cost and value functions give rise to an additive 6-spanner. In Section 2.1, we show that different costs and values lead to other sparseness-distortion tradeoffs, including a new additive 2-spanner of size $O(n^{3/2})$ that is quite different from previous constructions [Aingworth et al. 1999; Dor et al. 2000; Elkin and Peleg 2004; Thorup and Zwick 2006].

$$\begin{aligned} \text{value}(P) &= |\{\{C, C'\} \subseteq \mathcal{C}(P) : \delta_P(C, C') < \delta_H(C, C')\}|, \\ \text{cost}(P) &= |P \setminus H|. \end{aligned}$$

Note that the cost and value of a path is with respect to a subgraph H , which is our spanner under construction. The cost of a path is the number of its edges not already included in the spanner. The value function represents, roughly, how much the intercluster distances would be improved if P were included in the spanner.

Our path buying algorithm (Phase 2) is given in Figure 3. It refers to a set \mathcal{P} of shortest paths between all pairs of vertices with the following restrictions. If $P \in \mathcal{P}$ then all subpaths of P are also in \mathcal{P} . Furthermore, for every three consecutive vertices $\langle u_1, u_2, u_3 \rangle$ in a path $P \in \mathcal{P}$, if $\mathcal{C}(u_1) = \mathcal{C}(u_3)$, then u_2 is the center of $\mathcal{C}(u_1)$. This helps to reduce the cost of paths since $\langle u_1, u_2, u_3 \rangle \subseteq H_0$.

The remainder of the proof is structured as follows. In Lemma 2.3 we argue that, in the sum of values of paths purchased, the number of times any cluster pair is counted is bounded by a constant. This implies that the sum of values is $O(n^{4/3})$, since $|\mathcal{C}| = n^{2/3}$, and by our criterion for purchasing paths, that the sum of costs is also $O(n^{4/3})$. In Lemma 2.4 we relate the cost of a path to the number of clusters intersecting it. Finally, and most importantly, Lemma 2.5 shows that if any shortest path is too expensive to be purchased then the existing spanner edges already guarantee a path with additive distortion at most 6.

In the following lemmas, $\text{value}(P)$ and $\text{cost}(P)$ represent the value and cost of P at the time it was considered by the path buying algorithm in Figure 3.

LEMMA 2.3. *Let $H = H_0 \cup P_1 \cup P_2 \cup \dots \cup P_p$, where P_i is the i th path bought in the path buying phase. Then $\sum_{i=1}^p \text{value}(P_i) \leq 5 \binom{|\mathcal{C}|}{2} < \frac{5}{2} n^{4/3}$.*

PROOF. Let $H_i = H_0 \cup P_1 \cup \dots \cup P_i$. For any two clusters $C, C' \in \mathcal{C}$, let $P(C, C') = \{P_{j(1)}, P_{j(2)}, \dots, P_{j(r)}\}$ be those purchased paths such that

$$\delta_{P_{j(k)}}(C, C') < \delta_{H_{j(k)-1}}(C, C') \quad \text{for } k \in [1, r].$$

$H \leftarrow H_0$	{edges chosen in clustering}
For each shortest path $P \in \mathcal{P}$	
If $2 \cdot \text{value}(P) \geq \text{cost}(P)$	{if the path is a bargain}
then $H \leftarrow H \cup P$	{then buy it!}
Return H	

FIG. 3. The path buying algorithm. \mathcal{P} is a set of $\binom{n}{2}$ shortest paths between all pairs of vertices.

By the definition of the value function, $\sum_{i=1}^p \text{value}(P_i) = \sum_{\{C, C'\} \subseteq \mathcal{C}} |P(C, C')|$. Since $P_{j(1)}$ is a *shortest* path in G , we have that $\delta_{P_{j(1)}}(C, C') \leq \text{diam}_G(C) + \delta_G(C, C') + \text{diam}_G(C') \leq \delta_G(C, C') + 4$. This implies that $|P(C, C')| \leq 5$ since $\delta_G(C, C') \leq \delta_{P_{j(r)}}(C, C') < \delta_{P_{j(r-1)}}(C, C') < \dots < \delta_{P_{j(1)}}(C, C') \leq \delta_G(C, C') + 4$. That is, after $P_{j(1)}$ is purchased the distance from C to C' can only be improved four more times. \square

LEMMA 2.4. *If $P \in \mathcal{P}$ then either $|\mathcal{C}(P)| = 1$ or there exists a subpath $P' \subseteq P$ such that $\mathcal{C}(P') = \mathcal{C}(P)$ and $\text{cost}(P') \leq 2|\mathcal{C}(P')| - 3$.*

PROOF. Let $P = \langle u, \dots, u' \rangle$ and $P' \subseteq P$ be minimal such that $\mathcal{C}(P) = \mathcal{C}(P')$; see Figure 4. (This means that if the first or last cluster of P has two or three vertices in common with P then only the innermost one appears in P' .) The only edges in P' that might not be in $H \supseteq H_0$ are those between clustered vertices. Furthermore, if three consecutive vertices u_1, u_2, u_3 belong to the same cluster then u_2 is the center of cluster and $\langle u_1, u_2, u_3 \rangle \subseteq H_0$; see Figure 4. Thus the total number of intercluster edges and intracluster edges that are not in H are bounded by $|\mathcal{C}(P')| - 1$ and $|\mathcal{C}(P')| - 2$. \square

LEMMA 2.5. *The subgraph H returned by the path buying algorithm is happy.*

PROOF. Let $P = \langle u, \dots, u' \in \mathcal{P} \rangle$ be the shortest path from u to u' in G . By the statement of Property 2.1, we can dispense with several trivial cases and assume that P was not purchased in phase two, that both u and u' are clustered, and that $\mathcal{C}(u) \neq \mathcal{C}(u')$. Let $P' \subseteq P$ be the subpath guaranteed by Lemma 2.4. The case when P' is included in H is also trivial. Thus we have the following inequalities:

$$2 \cdot \text{value}(P') < \text{cost}(P') \leq 2 \cdot |\mathcal{C}(P')| - 3, \quad (1)$$

where the first inequality follows from the fact that P' was not included in H and the second from Lemma 2.4. Define A as the set of cluster pairs:

$$A = \left\{ \{C_0, C_1\} : \begin{array}{l} C_0 \in \{\mathcal{C}(u), \mathcal{C}(u')\}, C_1 \in \mathcal{C}(P') \setminus \{C_0\} \\ \text{and } \delta_{P'}(C_0, C_1) < \delta_H(C_0, C_1) \end{array} \right\}.$$

The cluster pairs counted in A are also counted in $\text{value}(P')$ so $|A| \leq \text{value}(P')$. By the inequalities of Equation (1), $\text{value}(P') \leq |\mathcal{C}(P')| - 2$. Notice that the maximum number of cluster pairs that could be counted by A is $2|\mathcal{C}(P')| - 3$.⁷ This means that, for at least $|\mathcal{C}(P')| - 1$ of these cluster pairs, their distance in the spanner is no worse than their distance in P' . By the pigeonhole principle, there must be some

⁷There is the pair $\{\mathcal{C}(u), \mathcal{C}(u')\}$ and the $2(|\mathcal{C}(P')| - 2)$ pairs $\{C_0, C_1\}$, where $C_0 \in \{\mathcal{C}(u), \mathcal{C}(u')\}$ and $C_1 \in \mathcal{C}(P') \setminus \{\mathcal{C}(u), \mathcal{C}(u')\}$.




FIG. 4. If the first (and last) cluster intersects P at more than one vertex, all but the innermost one can be left out of P' .

cluster $C^* \in \mathcal{C}(P') = \mathcal{C}(P)$ satisfying both

$$\delta_H(\mathcal{C}(u), C^*) \leq \delta_{P'}(\mathcal{C}(u'), C^*)$$

and

$$\delta_H(C^*, \mathcal{C}(u')) \leq \delta_{P'}(C^*, \mathcal{C}(u')).$$

Since $\mathcal{C}(P) = \mathcal{C}(P')$, it also follows that P has this property. \square

LEMMA 2.6. *If H is the subgraph purchased by the path buying algorithm then $|H| < 6 \cdot n^{4/3}$.*

PROOF. One can easily see that $|H| = |H_0| + \sum_i \text{cost}(P_i)$. By construction we have $|H_0| \leq n^{4/3}$. It follows from Lemma 2.3 that $\sum_i \text{cost}(P_i) \leq 2 \cdot \sum_i \text{value}(P_i) < 5 \cdot n^{4/3}$. \square

THEOREM 2.7. *There exists an additive 6-spanner of any graph with size $O(n^{4/3})$.*

PROOF. Follows from Lemmas 2.2, 2.5, and 2.6. \square

2.1. PARAMETERIZING THE PATH BUYING ALGORITHM. Our 6-spanner construction can be generalized to larger additive distortion. However, we can only guarantee that a β -spanner ($\beta > 6$) has $o(n^{4/3})$ edges if the graph satisfies additional requirements. Let $\Gamma_k(G)$ be the number of edges in G that lie on some cycle with length at most $2k$. (Notice that, when $k = 1$, $\Gamma_1(G) = 0$ since all graphs have girth at least 3.) Theorem 2.8 says that when the graph is not *too* far from having girth greater than $2k$ there exist a number of additive spanners with different size-distortion tradeoffs, the sparsest of which has size $\Theta(n^{1+1/(2k+1)})$. Figure 5 lists some of the corollaries of Theorem 2.8.

THEOREM 2.8. *For any graph G and integer parameters $k \geq 1$ and $\ell \in [0, k]$, there exists an additive $(2k + 4\ell)$ -spanner with size $O(\Gamma_k(G) + n^{1+\frac{1}{k+\ell+1}})$. Furthermore, the spanner can be constructed in time $O(mn^{1-\frac{\ell}{k+\ell+1}})$.*

For $k > 1$ Theorem 2.8 is quite weak since there is no reason to believe that $\Gamma_k(G)$ is small. We expressed the size bound in Theorem 2.8 in terms of $\Gamma_k(G)$ not because high-girth graphs are particularly interesting but because they serve to highlight the difficulties in generalizing our approach. The most natural way to extend the path-buying algorithm of Section 2 is to partition the graph into radius-2 clusters and define some kind of value function that takes the distances between these clusters into account. This approach works out very well if either (1) every vertex in the graph lies in some cluster or (2) the graph has girth greater than 4. It

Add. Dist.	Spanner Size	Construction Time	Notes
2	$O(n^{3/2})$	$O(mn)$	$k = 1, \ell = 0$
6	$O(n^{4/3})$	$O(mn^{2/3})$	$k = \ell = 1$
GRAPHS WITH GIRTH > 4			
0	$O(n^{3/2})$	—	Wenger [1991]
4	$O(n^{4/3})$	$O(mn)$	$k = 2, \ell = 0$
8	$O(n^{5/4})$	$O(mn^{3/4})$	$k = 2, \ell = 1$
12	$O(n^{6/5})$	$O(mn^{3/5})$	$k = \ell = 2$
GRAPHS WITH GIRTH > 6			
0	$O(n^{4/3})$	—	Wenger [1991]
6	$O(n^{5/4})$	$O(mn)$	$k = 3, \ell = 0$
10	$O(n^{6/5})$	$O(mn^{4/5})$	$k = 3, \ell = 1$
14	$O(n^{7/6})$	$O(mn^{2/3})$	$k = 3, \ell = 2$
18	$O(n^{8/7})$	$O(mn^{4/7})$	$k = \ell = 3$

FIG. 5. Additive spanners derived from Theorem 2.8. See Wenger [1991] for constructions of graphs with girth > 4 and > 6 , and Thorup and Zwick [2005] for references to other constructions of high-girth graphs. See Aingworth et al. [1999], Dor et al. [2000], Elkin and Peleg [2004], and Thorup and Zwick [2006] for other constructions of additive 2-spanners.

is not possible to guarantee (1) without creating too many clusters and assuming (2) is unrealistic. Thus, dealing with 3- and 4-cycles is, strangely, *the* obstacle to generalizing the path-buying algorithm.

The spanners provided by Theorem 2.8 are constructed with a path buying algorithm very similar to the one in the Section 2. In our additive 6-spanner construction we guaranteed that if a shortest path P from u to u' is *not* purchased then there is a path in the spanner between $\mathcal{C}(u)$ and $\mathcal{C}(u')$ via a special cluster C^* . The additive error of 6 reflects the cost of traversing three radius-1 clusters. The parameterized path-buying algorithm in this section makes the same type of guarantee, except that C^* will have radius k and $\mathcal{C}(u)$ and $\mathcal{C}(u')$ will have radius ℓ , where $0 \leq \ell \leq k$. Roughly speaking, the parameterized value function will count pairs of clusters with radii ℓ and k .

Before running the path buying algorithm, we compute an appropriate clustering and an initial set of edges $H_{k,\ell}$, which plays the same role as H_0 in our 6-spanner. This clustering procedure that produces $H_{k,\ell}$ is very similar to the randomized clustering procedure presented later in Section 3, though the analysis of these two schemes is sufficiently different to justify two expositions. We begin by selecting vertex sets V_ℓ and V_k where V_ℓ is a random sample of V of size $n^{1-\ell\epsilon}$ and V_k is a random sample of V_ℓ of size $n^{1-k\epsilon}$, where ϵ will be chosen later. We find two clusterings $\mathcal{C}_\ell, \mathcal{C}_k$, where, for $i \in \{\ell, k\}$, \mathcal{C}_i consists of a set of disjoint subsets of V . Each $C \in \mathcal{C}_i$ is centered at some vertex in V_i and a vertex v is contained in some $C \in \mathcal{C}_i$ if and only if $\delta(v, V_i) \leq i$; in particular, the distance from v to the center of its cluster is at most i . $H_{k,\ell}$ consists of a radius- i spanning tree of each $C \in \mathcal{C}_i$ and $i \in \{k, \ell\}$, as well as *every* edge incident to a vertex that does not appear in *both* clusterings \mathcal{C}_ℓ and \mathcal{C}_k . The number of edges contributed by the spanning trees is only $2n$, so we are mainly concerned with the expected number of the remaining edges.

LEMMA 2.9. *Given a graph G with m edges, the subgraph $H_{k,\ell}$ can be constructed in $O(m)$ time and $\mathbb{E}[|H_{k,\ell}|] = O(\Gamma_k(G) + n^{1+\epsilon})$.*

PROOF. To simplify the analysis, we imagine constructing $H_{k,\ell}$ by selecting vertex sets $V = V_0 \supseteq V_1 \supseteq \dots \supseteq V_k$, where V_i is a random sample of V_{i-1} of size $n^{1-i\epsilon}$. We construct the clusterings $\mathcal{C}_0, \dots, \mathcal{C}_k$ iteratively. Let v be a vertex incident to clusters $C_1, C_2, \dots, C_s \in \mathcal{C}_i$. If any of the centers of C_1, \dots, C_s is sampled for inclusion in V_{i+1} , v will join one adjacent sampled cluster that will be included in \mathcal{C}_{i+1} . When none of the clusters C_1, \dots, C_s are sampled, we include *all* of v 's incident edges in $H_{k,\ell}$. We now consider the expected number of edges contributed by v . The probability that v is included in \mathcal{C}_{i+1} is at least⁸ $1 - (1 - n^{-\epsilon})^s$. The effect of v *not* appearing in \mathcal{C}_{i+1} is to include in $H_{k,\ell}$ all edges incident to v , which could be significantly larger than both s and n^ϵ . Let $C_1, \dots, C_{s'}$ be those clusters connected to v by at most one edge. Every edge (v, w) connecting v to a cluster C appearing in $\mathcal{C}_{s'+1}, \dots, \mathcal{C}_s$ lies on a cycle of length at most $2i \leq 2k$: the one consisting of (v, w) , a path from w to $w' \in C$ passing through the center of C , and the edge (w', v) . Such a $w' \neq w$ exists because v is connected to C by at least two edges. If v is unclustered in \mathcal{C}_{i+1} then the contribution of the edges from v to $\mathcal{C}_{s'+1}, \dots, \mathcal{C}_s$ is counted in the $\Gamma_k(G)$ term. Excluding these edges, the expected number of edges contributed by v is at most $(s' - 1) \cdot \Pr[v \text{ does not appear in } \mathcal{C}_{i+1}] \leq (s' - 1)(1 - n^{-\epsilon})^s < n^\epsilon$.⁹ \square

Before proving Theorem 2.8, we establish a number of small lemmas that closely parallel Lemmas 2.2, 2.3, 2.4, 2.5, and 2.6. Due to the complications of having clusters with different radii, we need to use a new value function, a new definition of *happiness*, and a slightly different version of the path buying algorithm.

Whereas the previous definition of happiness and the previous value function considered pairs of radius-1 clusters, the new definitions will consider pairs of the form (v, C) , where $v \in V_\ell$ is the *center* of a radius- ℓ cluster and $C \in \mathcal{C}_k$ is a radius- k cluster.

PROPERTY 2.10. *A subgraph H of G that contains $H_{k,\ell}$ is happy if, for any $u, u' \in V_\ell$, there is a shortest path P (with respect to G) from u to u' and a cluster $C^* \in \mathcal{C}_k(P)$ such that, $\delta_H(u, C^*) \leq \delta_P(u, C^*)$ and $\delta_H(u', C^*) \leq \delta_P(u', C^*)$.*

As in Lemma 2.2, Lemma 2.11 reduces the problem to finding a happy subgraph.

LEMMA 2.11. *Any happy subgraph of G is an additive $(2k + 4\ell)$ -spanner of G .*

PROOF. Consider a shortest path between any two vertices v, v' that are clustered in both \mathcal{C}_ℓ and \mathcal{C}_k and let u and u' be the centers of the \mathcal{C}_ℓ clusters containing v and v' , respectively. See Figure 6. Let P be the shortest path between u and u' . By the happiness of H , there exists a cluster $C^* \in \mathcal{C}_k(P)$ such that

$$\delta_H(u, u') \leq \delta_H(u, C^*) + \text{diam}(C^*) + \delta_H(C^*, u') \leq \delta(u, u') + 2k.$$

⁸This would be the *exact* probability if V_{i+1} were selected from V_i by sampling each element with probability $n^{-\epsilon}$. Selecting V_{i+1} as a random subset of V_i with size $n^{1-(i+1)\epsilon}$ only improves v 's chance of being clustered in \mathcal{C}_{i+1} .

⁹For the last inequality, observe that $s'(1 - 1/x)^{s'} \leq x(s'/x)e^{-s'/x} \leq x$, for any positive x .




FIG. 6. The clusters $\mathcal{C}_l(u)$, $\mathcal{C}_l(u')$, and $C^* \in \mathcal{C}_k$, are indicated by ovals. Dashed curves represent paths in the spanner H . The shortest paths in G connecting v to v' and u to u' are indicated.

Using this bound on the distance from u to u' in H , we can bound $\delta_H(v, v')$ as

$$\begin{aligned} \delta_H(v, v') &\leq \delta_H(v, u) + \delta_H(v', u') + \delta_H(u, u') \\ &\leq 2\ell + \delta(u, u') + 2k \\ &\leq \delta(v, v') + 2k + 4\ell. \quad \square \end{aligned}$$

Let \mathcal{P} be a set of shortest paths between all pairs of vertices. We only consider shortest paths in \mathcal{P} and insist that it satisfy two properties. First, any subpath of a path in \mathcal{P} is also in \mathcal{P} . Second, if $\langle u_0, u_1, \dots, u_{2k} \rangle \in \mathcal{P}$ and $\mathcal{C}_k(u_0) = \mathcal{C}_k(u_{2k})$, then u_k is the center of $\mathcal{C}_k(u_0)$ and $\langle u_0, \dots, u_{2k} \rangle$ is contained in the spanning tree of $\mathcal{C}_k(u_0)$ included in $H_{k,\ell}$. Let $P(u, u') \in \mathcal{P}$ be the shortest path between u and u' .

We use the same cost function as before: $\text{cost}(P) = |P \setminus H|$, where H now represents the spanner under construction. However, the value function is *only* defined on paths connecting vertices in V_ℓ . Letting $P = P(u, u')$, where $u, u' \in V_\ell$, we define $\text{value}(P)$ as

$$\text{value}(P) = |\{(w, C) : w \in \{u, u'\}, C \in \mathcal{C}_k(P), \text{ and } \delta_P(w, C) < \delta_H(w, C)\}|.$$

The path buying algorithm below has a few differences from the 6-spanner algorithm. We only consider shortest paths between vertices in V_ℓ and the criterion for purchasing a path P is slightly stronger: $2k \cdot \text{value}(P)$ (rather than $2 \cdot \text{value}(P)$) must be at least $\text{cost}(P)$.

LEMMA 2.12. *For any $k \geq 1$ and $\ell \in [0, k]$, the size of the spanner returned by the path buying algorithm is $O(\Gamma_k(G) + n^{1+1/(k+\ell+1)})$.*

PROOF. Let P_1, \dots, P_p be the paths purchased. It follows that the size of H is exactly $|H_{k,\ell}|$ plus

$$\sum_{1 \leq i \leq p} \text{cost}(P_i) \leq 2k \sum_{1 \leq i \leq p} \text{value}(P_i) \leq 2\mu k |V_\ell \times \mathcal{C}_k| = 2\mu k n^{2-(k+\ell)\epsilon},$$

where μ is the maximum number of times that any pair (v, C) , where $v \in V_\ell$, $C \in \mathcal{C}_k$ could be counted in $\sum_i \text{value}(P_i)$. One can see that $\mu = 2k + 1$. When (v, C) is first counted, say when P_i is purchased, we have $\delta_{P_i}(v, C) \leq \delta(v, C) + \text{diam}(C) = \delta(v, C) + 2k$. After P_i is purchased the distance between v and C in H can only be improved $2k$ more times. Thus, the size of H is $O(\Gamma_k(G) + n^{1+\epsilon} + k^2 n^{2-(k+\ell)\epsilon})$. Setting $n^\epsilon = (k^2 n)^{1/(k+\ell+1)}$, the total size is $O(\Gamma_k(G) + n^{1+1/(k+\ell+1)})$. \square

$H \leftarrow H_{k,\ell}$	{edges chosen in clustering}
For each $\{u, u'\} \subseteq V_\ell$ let $P = P(u, u')$	
If $2k \cdot \text{value}(P) \geq \text{cost}(P)$	{if the path is a bargain}
then $H \leftarrow H \cup P$	{then buy it!}
Return H	

FIG. 7. Constructing an additive $(2k + 4\ell)$ -spanner.

It only remains to show that H is happy and, in particular, that if a path $P = P(u, u')$ is not purchased then $\delta_H(u, u') \leq \delta(u, u') + 2k$.

LEMMA 2.13. *The spanner returned by the path buying algorithm is happy.*

PROOF. We first bound $\text{cost}(P)$ then show that, if P is not purchased, there must be some $C^* \in \mathcal{C}_k(P)$ satisfying Property 2.10. Consider any cluster $C \in \mathcal{C}_k(P)$ and let z, z' be the first and last vertices in P that belong to C . An edge is internal to C if it lies on the subpath $\langle z, \dots, z' \rangle$. (Note that not all vertices between z and z' are necessarily in C .) It follows from our choice of shortest paths \mathcal{P} that either $|\langle z, \dots, z' \rangle| \leq 2k - 1$ or $\langle z, \dots, z' \rangle \subseteq H_{k,\ell} \subseteq H$. Thus, the total number of internal edges in $P \setminus H$ is at most $(2k - 1)|\mathcal{C}_k(P)|$ and the number of remaining edges in $P \setminus H$ is at most $|\mathcal{C}_k(P)| - 1$. In total we have $\text{cost}(P) \leq 2k|\mathcal{C}_k(P)| - 1$. If we fail to purchase P then $2k \cdot \text{value}(P) \leq \text{cost}(P) - 1 \leq 2k|\mathcal{C}_k(P)| - 2$, implying that $\text{value}(P) \leq |\mathcal{C}_k(P)| - 1$. Since there are $2|\mathcal{C}_k(P)|$ pairs of the form (w, C) where $w \in \{u, u'\}$ and $C \in \mathcal{C}_k(P)$, our bound on $\text{value}(P)$ implies that, for at least one $C^* \in \mathcal{C}_k(P)$, $\delta_H(u, C^*) \leq \delta_P(u, C^*)$ and $\delta_H(u', C^*) \leq \delta_P(u', C^*)$. \square

It follows directly from Lemmas 2.11, 2.12, and 2.13 that the path buying algorithm of Figure 7 returns an additive $(2k + 4\ell)$ -spanner H with the claimed size. However, any implementation of the Figure 7 algorithm, *as it is stated*, would be prohibitively slow. The most difficult problem is updating path costs and values after each shortest path is purchased. In the implementation described below, we use a different pair of cost and value functions that are easier to evaluate and update.

Rather than maintain the actual distance between a $v \in V_\ell$ and $C \in \mathcal{C}_k$, we keep an upper bound $\hat{\delta}_H(v, C)$, which is the minimum distance between v and C in a path already purchased. We ignore the cost function and consider the valid upper bound $\text{cost}(P) \leq 2k|\mathcal{C}_k(P)| - 1$. (These changes have no effect on the correctness of Lemmas 2.12 and 2.13.) At the beginning of the path buying phase we construct, in $O(m|\mathcal{C}_\ell|)$ time, a shortest-path tree originating at each vertex in V_ℓ . In such a tree a *relevant* vertex is one in V_ℓ or one that has at least two V_ℓ vertices in different subtrees, that is, a branching vertex in the subtree connecting V_ℓ nodes. For $u \in V_\ell$, let R_u be the set of relevant vertices in u 's shortest path tree. We consider, in nondecreasing order by length, the shortest paths between pairs (u, u') where $u \in V_\ell$ and $u' \in R_u$. For a path $P = \langle u_0, u_1, \dots, u_q \rangle$, where $u_0 \in V_\ell$, let v be defined as

$$v(u_0, u_q) = |\{(u_0, C) : C \in \mathcal{C}_k(P) \text{ and } \delta_P(u_0, C) < \hat{\delta}_H(u_0, C)\}|.$$

We assume that after P is considered we have computed $v(u_0, u_i)$ and $|\mathcal{C}_k(\langle u_0, \dots, u_i \rangle)|$, for all relevant vertices $u_i \in R_{u_0}$. When we consider an extension of P , say $P' = \langle u_0, \dots, u_q, \dots, u_{q+r} \rangle$, where u_{q+r} is the first relevant vertex on the extension, we can easily update $|\mathcal{C}_k(P')|$ and $v(u_0, u_{q+r})$ in just $O(k + r)$ time. If $u_{q+r} \in V_\ell$ then we need to decide whether to purchase P' . We

check if $v(u_0, u_{q+r}) + v(u_{q+r}, u_0) \geq |\mathcal{C}_k(P')|$, which serves the same purpose as checking the inequality $2k \cdot \text{value}(P') \geq \text{cost}(P')$ in the algorithm in Figure 7. If the inequality holds, we buy P' and update $v(u, u')$ where $u \in \{u_0, u_{q+r}\}$ and $u' \in R_u$. This amounts to checking, for each pair (u_0, C) counted in $v(u_0, u_{q+r})$ and each path $P'' = \langle u_0, \dots, u' \rangle$, whether $\delta_{P'}(u_0, C) \leq \delta_{P''}(u_0, C)$. If so, the pair (u_0, C) should no longer be counted in $v(u_0, u')$, if it was counted there at all. The total time for these updates is $O(k \cdot |\mathcal{C}_\ell|^2 \cdot |\mathcal{C}_k|) = O(n^{3 - \frac{2\ell+k}{k+\ell+1}})$ and the total time for performing breadth-first searches is $O(m \cdot |\mathcal{C}_\ell|) = O(mn^{1 - \frac{\ell}{k+\ell+1}})$. We can safely assume that $m \geq n^{1 + \frac{1}{k+\ell+1}}$; if not we would simply return the original graph as a trivial additive 0-spanner. For m above this threshold, the two time bounds are both $O(mn^{1 - \frac{\ell}{k+\ell+1}})$. This concludes the proof of Theorem 2.8.

2.2. ADDITIVE SPANNERS WITH POLYNOMIAL DISTORTION. The first obstacle to improving Theorem 2.8 is dealing with graphs with lots of 3- and 4-cycles. It is strange that short cycles should impede the discovery of more additive spanners since high-girth graphs are the most difficult instances when optimizing for *multiplicative* distortion. The recent lower bounds of Woodruff [2006] provide some circumstantial evidence that short cycles really do complicate the problem. His hard instances are actually composed entirely of complete bipartite graphs, where each edge appears on a huge number of 4-cycles.

Theorem 2.8 generalizes the 6-spanner construction by considering clusterings with wider radii. Another direction for generalization is to consider more hops between clusters. Theorem 2.14 follows from a small adjustment to the 6-spanner algorithm.

THEOREM 2.14. *Every graph on n vertices contains an additive $O(n^{1-3\epsilon})$ -spanner with size $O(n^{1+\epsilon})$, for any constant $\epsilon \in (0, \frac{1}{3})$.*

PROOF. The algorithm is nearly identical to the 6-spanner construction. We find a radius-1 clustering \mathcal{C} containing $n^{1-\epsilon}$ clusters, and let H_0 contain all edges incident to at least one unclustered vertex and a spanning tree of each cluster. We run the path buying algorithm from Figure 3 using the same cost function but a different value function:

$$\text{value}(P) = n^{-1+3\epsilon} \cdot |\{\{C, C'\} \subseteq \mathcal{C}(P) : \delta_P(C, C') < \delta_H(C, C')\}|.$$

The rationale for this value function is simple. If there are $n^{1-\epsilon}$ clusters and our desired spanner size is $O(n^{1+\epsilon})$, each cluster pair can only afford to be charged the cost of buying $O(n^{-1+3\epsilon})$ edges, that is, a small fraction of a single edge. Let P be some shortest path and P' and P'' be the prefix and suffix of P containing $n^{1-3\epsilon}$ distinct clusters. If P is not purchased then, by the same reasoning used before, there must be three not necessarily distinct clusters C', C^*, C'' where $C' \in \mathcal{C}(P')$, $C'' \in \mathcal{C}(P'')$, and $C^* \in \mathcal{C}(P)$ such that $\delta_H(C', C^*) \leq \delta_P(C', C^*)$ and $\delta_H(C'', C^*) \leq \delta_P(C'', C^*)$. That is, the subpath of P connecting C' to C'' is approximated by H to within an additive distortion of 6. To get from the first vertex of P to C' and from C'' to the last vertex of P , we use any standard multiplicative $O(\epsilon^{-1})$ -spanner with size $O(n^{1+\epsilon})$. The total additive distortion is therefore $6 + O(\epsilon^{-1}(\text{cost}(P') + \text{cost}(P''))) = O(n^{1-3\epsilon})$. The size of H_0 and the multiplicative spanner is $O(n^{1+\epsilon})$ and the number of edges included by the path buying algorithm is $\sum_i \text{cost}(P_i) \leq 2 \sum_i \text{value}(P_i) = O(n^{-1+3\epsilon} \cdot \binom{|\mathcal{C}|}{2}) = O(n^{1+\epsilon})$. \square

Theorem 2.14 is not particularly impressive, but, again, it illustrates the flexible nature of the path buying approach. The previous best additive spanner with size $O(n^{1+\epsilon})$ had distortion roughly $O(n^{1-2\epsilon})$ [Bollobás et al. 2003] and was substantially more complicated to construct. Pettie [2007] has recently shown that every graph has an additive $\tilde{O}(n^{\frac{9}{16}-\frac{7\epsilon}{8}})$ -spanner with size $O(n^{1+\epsilon})$. Theorem 2.14 improves on this bound when $\epsilon > \frac{7}{34}$.

3. A Simple $(k, k - 1)$ -Spanner in Linear Time

In this section, we will first show how to construct a $(2k - 1, 0)$ -spanner of size $O(kn^{1+1/k})$ in $O(km)$ deterministic time. Then we will extend this method to construct a $(k, k - 1)$ spanner of size $O(kn^{1+1/k})$ in $O(km)$ deterministic time.

The input graph is $G = (V, E)$. A *cluster* is simply a set of vertices and a *clustering* is a set of disjoint clusters. A vertex is *(un)clustered* in a clustering \mathcal{C} if it appears (does not appear) in some cluster of \mathcal{C} . In a clustering \mathcal{C} , for any clustered vertex u , denote by $\mathcal{C}(u)$ the cluster of \mathcal{C} that contains u . For clusters C and C' , let $E(C, C') = (C \times C') \cap E(G)$ be the set of edges between C and C' . Let $E(v, C)$ be the set of edges between the vertex v and vertices in C . A vertex v is adjacent to a cluster C if $E(v, C) \neq \emptyset$. In a similar manner, two clusters C and C' are adjacent to each other if $E(C, C') \neq \emptyset$.

Our constructions in this section are based on a set of $k + 1$ clusterings, $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_k$, where $\mathcal{C}_0 = \{\{v\} : v \in V(G)\}$, $\mathcal{C}_k = \emptyset$, and $|\mathcal{C}_i| \leq n^{1-i/k}$. Below, we give two methods for constructing appropriate sequences of clusterings. The edge set of our $(2k - 1)$ -spanner S is defined by the following two rules.

- Rule R1.* For each cluster $C \in \mathcal{C}_i$, there exists a tree in S that spans C and has radius¹⁰ at most i .
- Rule R2.* For each vertex v that is unclustered in \mathcal{C}_i and each cluster $C \in \mathcal{C}_{i-1}$ adjacent to v , some edge from $E(v, C)$ appears in S .

The construction of Theorem 3.1 is slightly weaker than that of Halperin and Zwick [1996]; however it is the starting point for our $(k, k - 1)$ -spanner.

THEOREM 3.1. *A $(2k - 1)$ -spanner of size $O(kn^{1+1/k})$ can be constructed in $O(km)$ deterministic time.*

PROOF. We first prove that Rules R1 and R2 give a $(2k - 1)$ -spanner; we then prove the size and time bounds. Let (u, v) be an arbitrary edge in the original graph. If $\delta_S(u, v) \leq (2k - 1)\delta_G(u, v)$ then S is clearly a $(2k - 1)$ -spanner. Let ℓ be minimum such that either u or v was unclustered in \mathcal{C}_ℓ and without loss of generality let the unclustered vertex be u . By Rule R2 there must be an edge in S from u to $\mathcal{C}_{\ell-1}(v)$; call this edge (u, w) . By Rule R1 there must be a path in S from w to v of length at most $2(\ell - 1)$, twice the radius of $\mathcal{C}_{\ell-1}(v)$. Since $\ell \leq k$, it follows immediately that $\delta_S(u, v) \leq 2k - 1$.

Given the clustering \mathcal{C}_i we show how to compute \mathcal{C}_{i+1} such that the number of edges added to the spanner due to Rules R1 and R2 are at most n and $n^{1+1/k}$,

¹⁰Recall that maximum distance between any two vertices in a subgraph is at most twice the radius of that subgraph.

Initially $S = \emptyset$ and $\mathcal{C}_0 = \{\{v\} : v \in V(G)\}$
 For i from 1 to k
 —Let \mathcal{C}_i be sampled from \mathcal{C}_{i-1} with prob. $n^{-1/k}$ (If $i = k$ let $\mathcal{C}_k = \emptyset$)
 —For each vertex v which does not belong to any cluster in \mathcal{C}_i do (concurrently):
 (R1) If v is adjacent to some $C \in \mathcal{C}_i$, add v to C and add some edge of $E(v, C)$ to S .
 (R2) Otherwise, add to S some edge from $E(v, C)$, for each $C \in \mathcal{C}_{i-1}$ adjacent to v .
 Return the $(2k - 1)$ -spanner S

FIG. 8. A randomized $(2k - 1)$ -spanner construction.

respectively (This construction is a simplified version of one described by Elkin [2004].) Initially $\mathcal{C}_{i+1} = \emptyset$. We define the priority of a cluster $C \in \mathcal{C}_i$ to be the number of adjacent vertices that are unclustered with respect to \mathcal{C}_{i+1} . We repeatedly choose a cluster $C \in \mathcal{C}_i$ with maximum priority. If $\text{priority}(C) \geq n^{(i+1)/k}$, we add a new cluster to \mathcal{C}_{i+1} consisting of C and all unclustered vertices adjacent to C . (If C has radius i then the cluster added to \mathcal{C}_{i+1} clearly has radius $i + 1$.) It follows that $|\mathcal{C}_{i+1}| \leq n^{1-(i+1)/k}$ and that the number of edges included in the spanner due to Rule R2 is $\sum_{C \in \mathcal{C}_i} \text{priority}(C)$, which is at most $|\mathcal{C}_i| (n^{(i+1)/k} - 1) < n^{1+1/k}$. The number of edges added due to Rule R1 is at most the number of clustered vertices in \mathcal{C}_{i+1} , that is, at most n . The clustering \mathcal{C}_{i+1} can easily be generated in linear time using a priority queue consisting of n buckets. \square

The randomized construction of Baswana and Sen [2007] constructs the clusterings in an even simpler manner. Rather than carefully selecting clusters from \mathcal{C}_i for inclusion in \mathcal{C}_{i+1} , they randomly sample all clusters from \mathcal{C}_i with probability $n^{-1/k}$. The alternative $(2k - 1)$ -spanner construction based on this method is given in Figure 8. With this algorithm, the *expected* size of the spanner is $O(kn^{1+1/k})$.

We next improve our $(2k - 1)$ -spanner construction to obtain a $(k, k - 1)$ -spanner. All the edges that we put in the $(2k - 1)$ -spanner S are of the form $E(v, C)$. Now let us also add edges of the form $E(C, C')$, which will consist of edges from clusters of one clustering to clusters of another clustering.

—*Rule R3.* For each i with $0 \leq i \leq k - 1$ and for each pair of adjacent clusters C, C' with $C \in \mathcal{C}_i$ and $C' \in \mathcal{C}_{k-1-i}$, some edge from $E(C, C')$ appears in S .

—*Rule R4.* For each $i \geq \lceil k/2 \rceil$ and each pair of adjacent clusters C, C' with $C \in \mathcal{C}_i$ and $C' \in \mathcal{C}_{i-1}$, some edge from $E(C, C')$ appears in S .

The number of edges included due to Rule R3 is bounded by $n^{1-i/k} n^{1-(k-1-i)/k} = n^{1+1/k}$, for each i . Similarly, at most $n^{1-i/k} n^{1-(i-1)/k} = n^{2-2i/k+1/k}$ edges are included due to R4, which is at most $n^{1+1/k}$ since $i \geq \lceil k/2 \rceil$. Our entire $(k, k - 1)$ -spanner construction is given in Figure 9. It consists of just those edges included by Rules R1–R4.

Implementing Rules R3 and R4 takes linear time for any fixed i . Once it is proved that Rules R1–R4 yield a $(k, k - 1)$ -spanner, we can conclude with the following theorem.

THEOREM 3.2. *A $(k, k - 1)$ -spanner of size $O(kn^{1+1/k})$ can be computed in $O(km)$ deterministic time.*

We now show that S is a $(k, k - 1)$ -spanner. Let $t = \lfloor k/2 \rfloor$. We need some more definitions. Call a vertex i -(un)clustered if it appears (does not appear) in clustering

- (R1-2) Compute a $(2k - 1)$ -spanner S with our construction.
- (R3) Add to S one edge from $E(C, C')$ for each adjacent pair $C \in \mathcal{C}_i$ and $C' \in \mathcal{C}_{k-1-i}$, for i from 0 to $k - 1$.
- (R4) Add to S one edge from $E(C, C')$ for each adjacent pair $C \in \mathcal{C}_i$, and $C' \in \mathcal{C}_{i-1}$, for i from $\lceil k/2 \rceil$ to $k - 1$.

FIG. 9. A simple linear time algorithm for constructing a $(k, k - 1)$ -spanner.

\mathcal{C}_i . The *center* of a cluster $C \in \mathcal{C}_i$ is a vertex $c \in C$ such that the distance from c to any other vertex in C is at most i , the radius of C . If v is i -clustered, let $c_i(v)$ be the center of $\mathcal{C}_i(v)$. In analyzing the stretch of spanner S , we shall make extensive use of the following notations.

$$f_i(v) = \begin{cases} v & \text{if } v \text{ is } i\text{-unclustered,} \\ c_i(v) & \text{if } v \text{ is } i\text{-clustered;} \end{cases} \quad r_i(v) = \begin{cases} 0 & \text{if } v \text{ is } i\text{-unclustered,} \\ i & \text{if } v \text{ is } i\text{-clustered.} \end{cases}$$

It is easy to observe that $\delta_S(v, f_i(v)) \leq r_i(v)$.

$$\sigma(j) = \begin{cases} \sum_{i=0}^j 2i & : j \geq 0, \\ \sum_{i=0}^{|j|-1} 2i & : j < 0. \end{cases}$$

It is easy to observe that the equalities $\sigma(-j) = \sigma(j - 1)$ and $\sigma(j) - 2j = \sigma(j - 1)$ hold for all j .

We shall now state three simple Lemmas which follow from the rules R1–R4 used for constructing the spanner. Lemma 3.3 is based on R2 and is already proved as part of Theorem 3.1.

LEMMA 3.3. *If v is ℓ -unclustered for $\ell \leq k - 1$, then for each edge $(u, v) \in E$,*

$$\delta_S(f_{\ell-1}(u), f_\ell(v)) \leq 2\ell - 1 - r_{\ell-1}(u).$$

The following Lemma uses the rule R3 in a crucial manner.

LEMMA 3.4. *If v is ℓ -clustered, then the following inequality holds for each edge $(u, v) \in E$.*

$$\delta_S(f_{k-1-\ell}(u), f_\ell(v)) \leq 2k - 2\ell - 1 + r_\ell(v) - r_{k-1-\ell}(u).$$

PROOF. If u is $(k - 1 - \ell)$ -unclustered then it follows from rule R2 that there is a path of length at most $2k - 2\ell - 3$ in S between u and v . Since there is a path of length $\leq r_\ell$ from v to $c_\ell(v)$ in S , the inequality holds. The other case is when u is $(k - 1 - \ell)$ -clustered. By Rule R3, there is an edge in S between $\mathcal{C}_{k-1-\ell}(u)$ and $\mathcal{C}_\ell(v)$. This gives us a path of length k between $c_{k-1-\ell}(u)$ and $c_\ell(v)$. Also note that $r_\ell(v) - r_{k-1-\ell}(u) = -(k - 1 - 2\ell)$. Therefore, the inequality $\delta_S(f_{k-1-\ell}(u), f_\ell(v)) \leq 2k - 2\ell - 1 + r_\ell(v) - r_{k-1-\ell}(u)$ holds in this case as well. \square

The following Lemma uses the rule R4 in a crucial manner.

LEMMA 3.5. *For each edge $(u, v) \in E$ and integer $\ell \geq \lceil \frac{k}{2} \rceil$, the following inequality holds.*

$$\delta_S(f_{\ell-1}(u), f_\ell(v)) \leq 2\ell - 1 + r_\ell(v) - r_{\ell-1}(u).$$

PROOF. Depending upon whether u and v appear in clusterings $\mathcal{C}_{\ell-1}$ and \mathcal{C}_ℓ , respectively, we have four cases: both are unclustered, both are clustered, u is clustered but v is not, and vice versa. For the first case, that is, $u \in \mathcal{C}_{\ell-1}$ and $v \in \mathcal{C}_\ell$, it follows from rule R4 that there is a spanner path between $f_{\ell-1}(u)$ and $f_\ell(v)$ of length 2ℓ . This fact combined with the observation that $r_\ell(v) = \ell$ and $r_{\ell-1}(u) = \ell - 1$ imply the inequality. The proof of the remaining three cases follow from Lemma 3.3 in a straight forward manner. \square

We are now ready to prove that S is indeed a $(k, k - 1)$ -spanner. Our analysis will be based on Lemmas 3.3, 3.4, and 3.5. We first analyze the case when k is an odd positive integer.

Analysis of the stretch for odd value of k . For the case when k is odd, we shall prove the following Theorem whose simple corollary would imply that S is $(k, k - 1)$ -spanner. Note that $k = 2t + 1$ since k is odd.

THEOREM 3.6. *Let $\langle u = v_0, v_1, \dots, v_q = v \rangle$ be a path between any $u, v \in V$ in the given graph, and S be the spanner computed by our algorithm for an odd integer k . Then for each $0 \leq i \leq q$ and integer $\ell \in [0, 2t]$,*

$$\delta_S(v_0, f_\ell(v_i)) \leq ki + r_\ell(v_i) + \sigma(t - \ell).$$

PROOF. We shall prove the statement of the theorem by induction on $i \geq 0$.

Base Case. ($i = 0$). It is easy to observe that $\delta_S(v_0, f_\ell(v_0)) \leq r_\ell(v_0)$. Moreover, $\sigma(\cdot)$ is always nonnegative, and hence the base case is true.

Induction step. Let us suppose that the statement is true for all integers up to $i - 1$. We shall prove the validity of the statement for i .

—For $\ell > t$: Consider the path from v_0 to $f_\ell(v_i)$ formed by concatenating the two shortest paths: $v_0 \rightsquigarrow f_{\ell-1}(v_{i-1})$ and $f_{\ell-1}(v_{i-1}) \rightsquigarrow f_\ell(v_i)$ in the spanner. Using the induction hypothesis, it follows that the length of the shortest path $v_0 \rightsquigarrow f_{\ell-1}(v_{i-1})$ is $k(i - 1) + r_{\ell-1}(v_{i-1}) + \sigma(t - \ell + 1)$. Moreover, it follows from Lemma 3.5 that there is a path in the spanner from $f_{\ell-1}(v_{i-1})$ to $f_\ell(v_i)$ consisting of at most $2\ell - 1 + r_\ell(v_i) - r_{\ell-1}(v_{i-1})$ edges. Hence we can infer that

$$\begin{aligned} \delta_S(v_0, f_\ell(v_i)) &\leq k(i - 1) + r_\ell(v_i) + 2\ell - 1 + \sigma(t - \ell + 1) \\ &= ki + r_\ell(v_i) + 2(\ell - t - 1) + \sigma(t - \ell + 1) \quad \{\text{for } k = 2t + 1\} \\ &= ki + r_\ell(v_i) - 2(t - \ell + 1) + \sigma(t - \ell + 1) \\ &= ki + r_\ell(v_i) + \sigma(t - \ell). \end{aligned}$$

—For $\ell \leq t$: There are two cases here.

Case 1: v_i is ℓ -unclustered. In this case, consider the path formed by concatenating the shortest paths $v_0 \rightsquigarrow f_{\ell-1}(v_{i-1})$ and $f_{\ell-1}(v_{i-1}) \rightsquigarrow f_\ell(v_i)$ in the spanner. Using the induction hypothesis, it follows that the length of the path $v_0 \rightsquigarrow f_{\ell-1}(v_{i-1})$ is at most $k(i - 1) + r_{\ell-1}(v_{i-1}) + \sigma(t - \ell + 1)$. Moreover, since v_i is ℓ -unclustered, it follows from Lemma 3.3 that the path $f_{\ell-1}(v_{i-1}) \rightsquigarrow v_i$ is of length at most $2\ell - 1 - r_{\ell-1}(v_{i-1})$. Hence, we can infer that

$$\begin{aligned} \delta_S(v_0, f_\ell(v_i)) &\leq k(i - 1) + r_\ell(v_i) + 2\ell - 1 + \sigma(t - \ell + 1) \quad \{\text{note } r_\ell(v_i) = 0\} \\ &= ki + r_\ell(v_i) - 2(t - \ell + 1) + \sigma(t - \ell + 1) \quad \{\text{for } k = 2t + 1\} \\ &= ki + r_\ell(v_i) + \sigma(t - \ell). \end{aligned}$$

Case 2: v_i is ℓ -clustered. In this case, Lemma 3.4 implies that there is a path $f_{k-1-\ell}(v_{i-1}) \rightsquigarrow f_\ell(v_i)$ in the spanner with length at most $2k - 2\ell - 1 + r_\ell(v_i) - r_{k-1-\ell}(v_{i-1})$. Moreover, using the induction hypothesis the shortest path from v_0 to $f_{k-1-\ell}(v_{i-1})$ in the spanner is of length $k(i-1) + r_{k-1-\ell}(v_{i-1}) + \sigma(t-k+\ell+1)$. Hence we can infer that

$$\begin{aligned} \delta_S(v_0, f_\ell(v_i)) &\leq ki + r_\ell(v_i) + k - 1 - 2\ell + \sigma(-(t-\ell)) \quad \{\text{for } k = 2t + 1\} \\ &= ki + r_\ell(v_i) + 2(t-\ell) + \sigma(-(t-\ell)) \\ &= ki + r_\ell(v_i) + 2(t-\ell) + \sigma(t-\ell-1) \quad \{\sigma(-j) = \sigma(j-1)\} \\ &= ki + r_\ell(v_i) + \sigma(t-\ell). \quad \square \end{aligned}$$

COROLLARY 3.7. *The spanner computed by our algorithm is a $(k, k-1)$ -spanner for any odd value of k .*

PROOF. Let $\langle u = v_0, v_1, \dots, v_q = v \rangle$ be any path of length q between any two vertices $u, v \in V$ in the graph. If v is t -clustered then it follows from Theorem 3.6 that there is a path of length $kq + t$ from u to $c_t(v)$. Together with the path from $c_t(v)$ to v of length at most t , we have a path of length $kq + 2t = kq + k - 1$ from u to v in S . If v is t -unclustered then $f_t(v) = v$ and $r_t(v) = 0$, implying a path of length kq from u to v . Hence the spanner S is indeed a $(k, k-1)$ -spanner for any odd value of k . \square

Analysis of the stretch for even value of k . We shall now prove the following theorem for the case when k is even. Note that $k = 2t$ since k is even.

THEOREM 3.8. *Let $\langle u = v_0, v_1, \dots, v_q = v \rangle$ be a path between any $u, v \in V$ in the given graph, and S be the spanner computed by our algorithm for an even integer k . Then for each $0 \leq i \leq q$ and integer $\ell \in [0, 2t-1]$,*

$$\delta_S(v_0, f_\ell(v_i)) \leq ki + r_\ell(v_i) + \sigma(t-\ell) + \ell - t - (i+t-\ell) \bmod 2.$$

PROOF. We shall prove the statement of theorem by induction on $i \geq 0$.
Base Case. ($i = 0$). It is easy to observe that $\delta_S(v_0, f_\ell(v_0)) \leq r_\ell(v_0)$. Moreover, $\sigma(j) - j - (j \bmod 2)$ is always nonnegative, and hence the base case is true.
Induction step. Let us suppose that the statement is true for all integers up to $i-1$. We shall prove the validity of the statement for i .

—*For $\ell \geq t$:* Consider the path from v_0 to $f_\ell(v_i)$ formed by concatenating the two shortest paths: $v_0 \rightsquigarrow f_{\ell-1}(v_{i-1})$ and $f_{\ell-1}(v_{i-1}) \rightsquigarrow f_\ell(v_i)$ in the spanner. It follows from the induction hypothesis that the shortest path from v_0 to $f_{\ell-1}(v_{i-1})$ in the spanner is of length $k(i-1) + r_{\ell-1}(v_{i-1}) + \sigma(t-\ell+1) + \ell - 1 - t - (i+t-\ell-2) \bmod 2$. Moreover using Lemma 3.5, there is a path in the spanner from $f_{\ell-1}(v_{i-1})$ to $f_\ell(v_i)$ of length at most $2\ell - 1 + r_\ell(v_i) - r_{\ell-1}(v_{i-1})$. Hence we can infer that

$$\begin{aligned} \delta_S(v_0, f_\ell(v_i)) &\leq k(i-1) + r_\ell(v_i) + \sigma(t-\ell+1) + \ell - 1 - t \\ &\quad - (i+t-\ell) \bmod 2 + 2\ell - 1 \\ &\leq ki + r_\ell(v_i) + \sigma(t-\ell+1) - 2(t-\ell+1) + \ell \\ &\quad - t - (i+t-\ell) \bmod 2 \\ &= ki + r_\ell(v_i) + \sigma(t-\ell) + \ell - t - (i+t-\ell) \bmod 2. \end{aligned}$$

—*For $\ell < t$:* There are two cases here.

Case 1: v_i is ℓ -unclustered. In this case, consider the path formed by concatenating the shortest paths $v_0 \rightsquigarrow f_{\ell-1}(v_{i-1})$ and $f_{\ell-1}(v_{i-1}) \rightsquigarrow f_\ell(v_i)$ in the spanner. Using the induction hypothesis, it follows that there is a path in the spanner from v_0 to $f_{\ell-1}(v_{i-1})$ of length at most $k(i-1) + r_{\ell-1}(v_{i-1}) + \sigma(t - \ell + 1) + \ell - 1 - t - (i + t - \ell - 2) \bmod 2$. Moreover, since v_i is ℓ -unclustered, it follows from Lemma 3.3 that the path $f_{\ell-1}(v_{i-1}) \rightsquigarrow v_i$ is of length at most $2\ell - 1 + r_\ell(v_i) - r_{\ell-1}(v_{i-1})$. Hence we can infer that

$$\begin{aligned} \delta_S(v_0, f_\ell(v_i)) &\leq k(i-1) + r_\ell(v_i) + \sigma(t - \ell + 1) + \ell - 1 - t + 2\ell \\ &\quad - 1 - (i + t - \ell) \bmod 2 \\ &\leq ki + r_\ell(v_i) + \sigma(t - \ell + 1) - 2(t - \ell + 1) + \ell - t \\ &\quad - (i + t - \ell) \bmod 2 \\ &= ki + r_\ell(v_i) + \sigma(t - \ell) + \ell - t - (i + t - \ell) \bmod 2. \end{aligned}$$

Case 2: v_i is ℓ -clustered. In this case, Lemma 3.4 implies that there is a path $f_{k-1-\ell}(v_{i-1}) \rightsquigarrow f_\ell(v_i)$ in the spanner with length at most $2k - 2\ell - 1 + r_\ell(v_i) - r_{k-1-\ell}(v_{i-1})$. Moreover, using the induction hypothesis the shortest path from v_0 to $f_{k-1-\ell}(v_{i-1})$ in the spanner is of length at most $k(i-1) + r_{k-1-\ell}(v_{i-1}) + \sigma(t - k + \ell + 1) + k - \ell - 1 - t - (i + t + \ell) \bmod 2$. Hence there is a path from v_0 to $f_\ell(v_i)$ in the spanner passing through $f_{k-1-\ell}(v_{i-1})$ with length

$$\begin{aligned} \delta_S(v_0, f_\ell(v_i)) &\leq k(i-1) + r_\ell(v_i) + \sigma(t - k + \ell + 1) + 3k - 3\ell - 2 \\ &\quad - t - (i + t + \ell) \bmod 2 \\ &= ki + r_\ell(v_i) + \sigma(-(t - \ell - 1)) + (4t - 4\ell - 2) + \ell \\ &\quad - t - (i + t - \ell) \bmod 2 \\ &= ki + r_\ell(v_i) + \sigma(t - \ell - 2) + (4t - 4\ell - 2) + \ell - t \\ &\quad - (i + t - \ell) \bmod 2 \\ &= ki + r_\ell(v_i) + \sigma(t - \ell) + \ell - t - (i + t - \ell) \bmod 2. \quad \square \end{aligned}$$

COROLLARY 3.9. *The spanner computed by our algorithm is a $(k, k-1)$ -spanner for any even value of k .*

PROOF. Let $\langle u = v_0, v_1, \dots, v_q = v \rangle$ be the shortest path between u and v in the given graph. There are two cases :

— *q is odd.* Consider the shortest path from $u = v_0$ to v in the spanner passing through $f_t(v)$. Theorem 3.8 implies that there is a path in the spanner S from u to $f_t(v)$ of length at most $kq + r_t(v) - 1$. Thus there is a path in S from u to v of length $kq + 2r_t(v) - 1$. Since $r_t(v) \leq t = k/2$, therefore,

$$\delta_S(u, v) \leq kq + 2 \cdot k/2 - 1 = kq + k - 1.$$

— *q is even.* Consider the shortest path from $u = v_0$ to v in the spanner passing through $f_{t-1}(v)$. Theorem 3.8 implies that there is a path in the spanner S from u to $f_{t-1}(v)$ of length at most $kq + r_{t-1}(v)$. Hence there is a path in S from u to v passing through $f_{t-1}(v)$ with length $kq + 2r_{t-1}(v) = kq + k - 2$. Thus

$$\delta_S(u, v) \leq kq + k - 2.$$

Hence the spanner is indeed a $(k, k-1)$ -spanner for even k too. \square

Theorems 3.6 and 3.8 together imply that S is a $(k, k - 1)$ -spanner for any positive integer k . Thus Theorem 3.2 is completely proved.

3.1. IMPLEMENTATION IN OTHER MODELS OF COMPUTATIONS. Our algorithm (Figure 9) for $(k, k - 1)$ -spanners can be adapted quite easily to other models of computation. The complexity of each of these algorithms is close to optimal under relevant measures.

- In the external memory model [Aggarwal and Vitter 1988] and the cache oblivious model [Frigo et al. 1999], a $(k, k - 1)$ -spanner of $O(kn^{1+1/k})$ size can be computed using the same number of I/O operations as that of sorting km items. Sorting is one of the most primitive tasks in both models and optimal algorithms are known.
- In the CRCW PRAM model [Karp and Ramachandran 1990], a $(k, k - 1)$ -spanner of expected size $O(kn^{1+1/k})$ can be computed in $O(k \log^* n)$ time and $O(km)$ work. The algorithm employs primitive parallel subroutines like computing the smallest element, semisorting, and multiset hashing.
- In the synchronous distributed model [Peleg 2000], a $(k, k - 1)$ -spanner of expected $O(kn^{1+1/k})$ size can be computed in $O(k)$ rounds with total message volume $O(k^2m)$.

In order to convey to the reader the ease with which the algorithm is adapted in these models, we provide below the execution of the $(k, k - 1)$ -spanner construction in distributed environment. Adaptation of the algorithm in the external-memory and CRCW PRAM environment as mentioned above is similar to the adaptation of $(2k - 1)$ -spanner algorithm of Baswana and Sen [2007] in these models.

3.2. DISTRIBUTED ALGORITHM FOR $(k, k - 1)$ -SPANNER. Rules R3 and R4 of our algorithm (Figure 9) can be executed in $O(k)$ rounds of communication in a distributed network, and, using the randomized algorithm from Figure 8, Rules R1 and R2 can also be executed in $O(k)$ rounds. The main result in this section is Theorem 3.10. Before we prove it, let us elaborate a little on our model of distributed computation. In a synchronized distributed network, the nodes of the network solve a problem by exchanging messages in discrete *rounds*. In each round one message may be sent across each link in each direction. We are interested in three measures: the number of rounds, the maximum length of any message sent (measured in units of $O(\log n)$ bits), and the total length of all messages sent. Clearly any protocol requiring R rounds, maximum message length L , and total volume V can be converted to one with parameters RL/U , U , and V , for any any $U \geq 1$. That is, Theorem 3.10 can be adapted to any synchronized network with a fixed maximum message length.

THEOREM 3.10. *In a synchronized distributed network G , a $(k, k - 1)$ -spanner of G with expected size $O(kn^{1+1/k})$ can be constructed in $O(k)$ rounds of communication. The total message volume is $O(k^2m)$ and the maximum message length is $O(n^{1/2+1/2k})$.*

PROOF. We compute the clusters $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_k = \emptyset$ using the randomized algorithm from Figure 8. Each vertex is the center of its cluster in $\mathcal{C}_0 = \{\{v\} : v \in V(G)\}$. With probability $n^{-1/k}$, each center in \mathcal{C}_i declares itself to also be a center in \mathcal{C}_{i+1} . These random choices are made *before* the first round of communication.

After \mathcal{C}_i is computed, every vertex tells its neighbors whether it is clustered in \mathcal{C}_i and, if it is, the identity of its center in \mathcal{C}_i and the highest $j \geq i$ for which that center is also a center in \mathcal{C}_j . For each vertex w that has a neighbor v already clustered in \mathcal{C}_{i+1} , w declares (w, v) to be a spanner edge (Rule R1) and declares its center in \mathcal{C}_{i+1} to be that of v . Every vertex w that did not become clustered in \mathcal{C}_{i+1} declares one edge from $E(w, C)$ to be in the spanner (Rule R2), for each $C \in \mathcal{C}_i$ adjacent to w . Rules R1 and R2 require $k - 1$ rounds of communication, plus one more to let clustered vertices in \mathcal{C}_{k-1} inform their neighbors of this fact. Each message sent so far has unit length.

Once $\mathcal{C}_0, \dots, \mathcal{C}_{k-1}$ are computed, we implement Rules R3 and R4. Consider Rule R3, a fixed $i \geq 0$, and a fixed cluster $C \in \mathcal{C}_{(k-1)/2-i}$.¹¹ Rule R1 has created a tree T of spanner edges rooted at the center of C . This tree is used to inform the center of C of all incident clusters in $\mathcal{C}_{(k-1)/2+i}$, and for each such cluster, one connecting edge. Once the center decides which edges to select for Rule R3 it broadcasts its choices back through T . The number of rounds for Rule R3 is clearly $O(k)$. The maximum necessary message length (for fixed i) is $|\mathcal{C}_{(k-1)/2+i}|$ since duplicate edges connecting the same clusters can be ignored. With high probability $|\mathcal{C}_{(k-1)/2+i}| = O(n^{1/2+1/2k-i/k})$. Summing over $i \geq 0$, the maximum message length is bounded by $O(n^{1/2+1/2k})$. For even k , i is always at least $1/2$, so in this case the maximum message length is $O(\sqrt{n})$. The total message volume for Rule R3 is $O(k^2m)$ since each edge can participate for $k/2$ values of i and for each, contribute $O(k)$ units of message volume. The implementation and analysis of Rule R4 is very similar to R3. \square

4. Conclusion

The main existential question in the field of spanners is whether, for any given size bound $O(n^{1+\epsilon})$, there exist additive $\beta(\epsilon)$ -spanners. In this article we introduced a general construction technique that might help to resolve the question of additive spanners for general graphs.

In Section 3 we addressed the problem of computationally efficient spanner constructions and gave partial answers to two problems of practical significance: what is the highest-quality spanner that can be constructed in linear time? and which spanners can be constructed distributively in $O(1)$ rounds? It seems implausible that any additive or $(1 + \epsilon, \beta)$ -spanners admit such efficient constructions.

One promising direction for future research is to develop approximate distance oracles for unweighted graphs whose distortion improves with distance. The ultimate goal would be to have oracles with constant additive distortion, though any improvement over the purely multiplicative distortion of Thorup and Zwick [2005], Baswana and Sen [2003], Mendel and Naor [2007], Baswana and Kavitha [2006], and Baswana et al. [2008] would be a start. For example, there is no known $(3 - \epsilon, \beta)$ -distance oracle with size $O(n^{3/2})$ whose query time is reasonably fast.

ACKNOWLEDGMENT. We would like to thank Uri Zwick, Mikkel Thorup, and Michael Elkin for many helpful comments.

¹¹ If k is odd then i is an integer. For even k , i is a half-integer.

REFERENCES

- ABRAHAM, I., BARTAL, Y., CHAN, H. T.-H., DHAMDHERE, K., GUPTA, A., KLEINBERG, J. M., NEIMAN, O., AND SLIVKINS, A. 2005. Metric embeddings with relaxed guarantees. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 83–100.
- ABRAHAM, I., BARTAL, Y., AND NEIMAN, O. 2006a. Advances in metric embedding theory. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*. 271–286.
- ABRAHAM, I., BARTAL, Y., AND NEIMAN, O. 2007. Embedding metrics into ultrametrics and graphs into spanning trees with constant average distortion. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 502–511.
- ABRAHAM, I., GAVOILLE, C., AND MALKHI, D. 2006b. On space-stretch trade-offs: upper bounds. In *Proceedings of the 18th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*. 217–224.
- AGGARWAL, A., AND VITTER, J. S. 1988. The input/output complexity of sorting and related problems. *Comm. ACM* 31, 9, 1116–1127.
- AINGORTH, D., CHEKURI, C., INDYK, P., AND MOTWANI, R. 1999. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.* 28, 4, 1167–1181.
- ALON, N., HOORY, S., AND LINIAL, N. 2002. The Moore bound for irregular graphs. *Graphs Combin.* 18, 1, 53–57.
- ALON, N., KARP, R. M., PELEG, D., AND WEST, D. B. 1995. A graph-theoretic game and its application to the k -server problem. *SIAM J. Comput.* 24, 1, 78–100.
- ALTHÖFER, I., DAS, G., DOBKIN, D., JOSEPH, D., AND SOARES, J. 1993. On sparse spanners of weighted graphs. *Discrete Comput. Geom.* 9, 81–100.
- BĀDOIU, M., INDYK, P., AND SIDIROPOULOS, A. 2007. Approximation algorithms for embedding general metrics into trees. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- BARTAL, Y. 1996. Probabilistic approximations of metric spaces and its algorithmic applications. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 184–193.
- BASWANA, S., GAUR, A., SEN, S., AND UPADHYAY, J. 2008. Distance oracles for unweighted graphs: Breaking the quadratic barrier with constant additive error. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP)*. Lecture Notes in Computer Science, vol. 5125. Springer, Berlin, Germany, 609–621.
- BASWANA, S., AND KAVITHA, T. 2006. Faster algorithms for approximate distance oracles and all-pairs small stretch paths. In *Proceedings of the 47th IEEE Symposium on Foundations of Computer Science (FOCS)*. 591–602.
- BASWANA, S., KAVITHA, T., MEHLHORN, K., AND PETTIE, S. 2005. New constructions of (α, β) -spanners and purely additive spanners. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 672–681.
- BASWANA, S., AND SEN, S. 2003. A simple linear time algorithm for computing a $(2k - 1)$ -spanner of $O(n^{1+1/k})$ size in weighted graphs. In *Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP)*.
- BASWANA, S., AND SEN, S. 2006. Approximate distance oracles for unweighted graphs in expected $O(n^2 \log n)$ time. *ACM Trans. Algor.* 2, 4, 557–577.
- BASWANA, S., AND SEN, S. 2007. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *J. Rand. Struct. Algor.* 30, 4, 532–563.
- BOLLOBÁS, B., COPPERSMITH, D., AND ELKIN, M. 2003. Sparse distance preservers and additive spanners. In *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 414–423.
- CHAN, H. T.-H., DINITZ, M., AND GUPTA, A. 2006. Spanners with slack. In *Proceedings of the 14th Annual European Symposium on Algorithms (ESA)*. 196–207.
- CHAN, H. T.-H., AND GUPTA, A. 2006. Small hop-diameter sparse spanners for doubling metrics. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 70–78.
- CHAN, H. T.-H., GUPTA, A., MAGGS, B. M., AND ZHOU, S. 2005. On hierarchical routing in doubling metrics. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 762–771.
- COHEN, E. 1998. Fast algorithms for constructing t -spanners and paths with stretch t . *SIAM J. Comput.* 28, 210–236.
- COHEN, E. 2000. Polylog-time and near-linear work approximation scheme for undirected shortest-paths. *J. ACM* 47, 132–166.

- COPPERSMITH, D., AND ELKIN, M. 2005. Sparse source-wise and pair-wise distance preservers. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 660–669.
- COWEN, L. J. 2001. Compact routing with minimum stretch. *J. Algor.* 28, 170–183.
- COWEN, L. J., AND WAGNER, C. G. 2004. Compact roundtrip routing in directed networks. *J. Algor.* 50, 1, 79–95.
- DOR, D., HALPERIN, S., AND ZWICK, U. 2000. All-pairs almost shortest paths. *SIAM J. Comput.* 29, 5, 1740–1759.
- ELKIN, M. 2004. Private communication.
- ELKIN, M. 2005. Computing almost shortest paths. *ACM Trans. Algor.* 1, 2, 283–323.
- ELKIN, M., EMEK, Y., SPIELMAN, D. A., AND TENG, S.-H. 2005. Lower-stretch spanning trees. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*. 494–503.
- ELKIN, M., AND PELEG, D. 2001. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*. 173–182.
- ELKIN, M., AND PELEG, D. 2004. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. *SIAM J. Comput.* 33, 3, 608–631.
- ELKIN, M., AND PELEG, D. 2005. Approximating k -spanner problems for $k \geq 2$. *Theoret. Comput. Sci.* 337, 1–3, 249–277.
- ELKIN, M., AND ZHANG, J. 2006. Efficient algorithms for constructing $(1 + \epsilon, \beta)$ -spanners in the distributed and streaming models. *Distrib. Comput.* 18, 5, 375–385.
- EMEK, Y., AND PELEG, D. 2004. Approximating minimum max-stretch spanning trees on unweighted graphs. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 261–270.
- ERDŐS P. 1963. Extremal problems in graph theory. In *Theory of Graphs and its Applications*. Publishing House of the Czechoslovak Academy of Sciences, Prague, Czechoslovakia, 29–36.
- FAKCHAROENPHOL, J., RAO, S., AND TALWAR, K. 2004. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.* 69, 3, 485–497.
- FRIGO, M., LEISERSON, C. E., PROKOP, H., AND RAMACHANDRAN, S. 1999. Cache-oblivious algorithms. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 285–298.
- GUDMUNDSSON, J., LEVCOPOULOS, C., AND NARASIMHAN, G. 2002. Fast greedy algorithms for constructing sparse geometric spanners. *SIAM J. Comput.* 31, 5, 1479–1500.
- HALPERIN, S., AND ZWICK, U. 1996. Unpublished.
- HAR-PELED, S., AND MENDEL, M. 2006. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.* 35, 5, 1148–1184.
- INDYK, P. 2001. Algorithmic aspects of low-distortion geometric embeddings. 42nd Annual IEEE Symposium on Foundations of Computer Science. Tutorial. <http://theory.lcs.mit.edu/~indyk/tut.html>.
- INDYK, P., AND MATOUŠEK, J. 2004. Low-distortion embedding of finite metric spaces. In *Handbook of Discrete and Computational Geometry*, 2nd ed. CRC Press, Boca Raton, FL.
- KARP, R. M., AND RAMACHANDRAN, V. 1990. Parallel algorithms for shared-memory machines. In *Handbook of Computer Science*. MIT Press, Cambridge, MA, 869–942.
- MENDEL, M., AND NAOR, A. 2007. Ramsey partitions and proximity data structures. *J. European Math. Soc* 9, 2, 253–275.
- NARASIMHAN, G., AND SMID, M. 2007. *Geometric Spanner Networks*. Cambridge University Press, Cambridge, U.K.
- PELEG, D. 2000. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia, PA.
- PELEG, D., AND SCHAFFER, A. A. 1989. Graph spanners. *J. Graph Theor.* 13, 99–116.
- PELEG, D., AND ULLMAN, J. D. 1989. An optimal synchronizer for the hypercube. *SIAM J. Comput.* 18, 740–747.
- PELEG, D., AND UPFAL, E. 1989. A trade-off between space and efficiency for routing tables. *J. ACM* 36, 3, 510–530.
- PETTIE, S. 2007. Low distortion spanners. In *Proceedings of the 34th International Colloquium on Automata, Languages, and Programming (ICALP)*. 78–89.
- RODITTY, L., THORUP, M., AND ZWICK, U. 2002. Roundtrip spanners and roundtrip routing in directed graphs. In *Proceedings of the 13th ACM-SIAM Symposium On Discrete Algorithms (SODA)*. 844–851.
- RODITTY, L., THORUP, M., AND ZWICK, U. 2005. Deterministic constructions of approximate distance oracles and spanners. In *Proceedings of the 32nd International Colloquium on Automata, Language, and Programming (ICALP)*. 261–272.

- RODITY, L., AND ZWICK, U. 2004. On dynamic shortest paths problems. In *Proceedings of the 12th Annual European Symposium on Algorithms (ESA)*. 580–591.
- SPIELMAN, D. A., AND TENG, S.-H. 2004. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*. 81–90.
- THORUP, M., AND ZWICK, U. 2001. Compact routing schemes. In *Proceedings of the 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*. 1–10.
- THORUP, M., AND ZWICK, U. 2005. Approximate distance oracles. *J. ACM* 52, 1, 1–24.
- THORUP, M., AND ZWICK, U. 2006. Spanners and emulators with sublinear distance errors. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 802–809.
- WENGER, R. 1991. Extremal graphs with no C^4 s, C^6 s, or C^{10} s. *J. Comb. Theor. Ser. B* 52, 1, 113–116.
- WOODRUFF, D. 2006. Lower bounds for additive spanners, emulators, and more. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 389–398.
- XIAO, J., LIU, L., XIA, L., AND JIANG, T. 2007. Fast elimination of redundant linear equations and reconstruction of recombination-free Mendelian inheritance on a pedigree. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 655–664.

RECEIVED JANUARY 2008; REVISED JULY 2008; ACCEPTED AUGUST 2008