Approximation Algorithms for Metric Embedding Problems

Kedar Dhamdhere
June 2005

School of Computer Science
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
R. Ravi, Chair
Anupam Gupta
Guy Blelloch
Piotr Indyk, MIT

Date and Location:
June 24, 2005
2 p.m.
Wean 4623

Submitted in partial fulfillment of the requirements
for the Degree of Doctor of Philosophy

This research was sponsored by the National Science Foundation (NSF) under grant nos. NSF CCF 04-30751 and ITR grant CCR-0122581 (The ALADDIN Project). The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the NSF, the U.S. government or any other entity.

Keywords: Approximation Algorithms, Metric Embedding, Distortion, Line metric, Spanning Trees

Abstract

We initiate the study of metric embedding problems from approximation point of view. Metric embedding is a map from the guest metric to the host metric. The quality of the embedding is defined in terms of distortion, the ratio by which pairwise distances get skewed in the host metric. While metric embeddings in general have received quite a lot of attention in theory community, most of the results about distortion prove uniform bounds that work for various families of host and guest metric. In this dissertation, we address the question: how to find the best embedding of the particular input metric into the host metric. We consider the real line as the host metric in our study.

We consider the following measures of quality of an embedding: distortion, average distortion and additive distortion. The distortion is the maximum ratio by which a pairwise distance gets stretched in a non-contracting embedding. We give $O(\sqrt{n})$-approximation for the distortion of embedding an unweighted graph metric to the metric. The average distortion is the ratio of average distance in the embedded metric to that in the input metric. We give a 17-approximation for the average distortion when embedding an arbitrary finite metric to a line metric. The additive distortion is the total absolute difference between input and output distances. We provide an $O(\sqrt{\log n})$-approximation for this objective function. We also show NP-hardness of these problems.

We also consider the problem of linear ordering of a metric, i.e. assigning numbers from 1 through n to the points in the metric, so as to minimize the 'stretch'. The stretch is the maximum pairwise distance in the ordering divided by the distance in the input metric. For this problem, we give $O(\log^3 n)$ approximation.

Finally, we consider the problem of constructing a probabilistic embedding of a graph into its spanning trees. We give a simple $O(\log^2 n)$-approximation algorithm that improves on the algorithm of Elkin et al. [27].

Introduction

Over the past decade, metric embeddings have been objects of much attention in theoretical computer science. This has been largely due to their many algorithmic applications, which range from simplifying the structure of the input data for approximation and online problems [6, 5, 12, 14, 30, 34, 40], serving as convenient relaxations of important NP-hard problems [8, 16, 17, 18, 33, 44] or simply by being the object of study [2, 31] arising from applications such as computational biology. Embedding techniques have become an indispensable addition to the algorithm designer's toolbox, providing powerful and elegant solutions to many algorithmic problems (see, e.g., [45, Chapter 15] and [38] for surveys).

An embedding of a metric (V, d) into *host* metric (H, δ) is a map $f : V \to H$. The quality of this map is measured by how closely the distances between points in d closely resemble those between their images in δ. An embedding f is called *non-contracting* if the map f does not decrease any of the distances, i.e., $d(x, y) \leq \delta(f(x), f(y))$ for all $x, y \in V$. (In the sequel, we will abbreviate $\delta(f(x), f(y))$ to $\delta(x, y)$.) An important measure of the quality of a non-contracting embedding f is the *distortion* $D = D(f)$, which is:

$$\text{distortion } D = \max_{x,y \in V} \frac{\delta(x, y)}{d(x, y)}. \tag{1}$$

(We note that a more general definition of distortion can be given that is scale-free; hence the restriction of non-contracting embeddings used here is without loss of generality.)

While many embedding techniques and algorithms are known, the analyses for these embeddings usually only offer uniform bounds on the distortion of the embeddings; few results which address the problem of minimizing the distortion required for embedding a *given* metric into the host space. In fact, very few results show how to even *approximate* the distortion to better than these uniform bounds.

This is perhaps best shown by a concrete example: Matoušek [46] proved that *any* metric (V, d) can be embedded into the real line with distortion $O(|V|)$; furthermore, the result is existentially tight, as the n-cycle cannot be embedded into the line with distortion $o(|V|)$ (see, e.g., [49, 36]). However, no algorithm is known for this problem which offers *per-instance* guarantees; even if a metric (X, d) may be embeddable into $\mathbb{R}$ with distortion $D = O(1)$, the known algorithms do not seem to guarantee that the embedding they output has distortion, say, that is within $O(|V|^{1-\epsilon})$ times D.

For most problems in metric embeddings: while uniform upper bounds are known for embeddings of many different families of metrics (e.g., general metrics, planar graph metrics, tree metrics) into a variety of host spaces (e.g., the Minkowski ℓ_p spaces, distributions of trees), very little is known about how to approximate the optimal distortion given a *fixed* metric (V, d) and a host space. One notable exception is the remark of Linial et al. [44] that the optimal embedding of any finite metric into (unbounded dimensional) Euclidean spaces to minimize distortion can be computed as a solution to a semi-definite program.

In this dissertation, we focus on studying the metric embeddings from an algorithmic perspective. In other words, we would like to address the questions of the form: 'given an input metric, how to best embed it into a prescribed host metric'? In particular, we focus on the line metric as the prescribed host metric. It turns out that many of these problems are NP-hard. Therefore, we look for approximation algorithms.

While distortion as defined above has been very popular, we also investigate other notions of the quality of the embedding. We describe each of these in detail in the following sections.

Let (V, d) denote the finite guest metric. We want to embed (V, d) into the line metric: $(\mathbb{R}, \delta)$. Let $|V| = n$ and let Δ be the diameter of the metric.

Average distortion

First we focus our attention on the average distortion of the embeddings arbitrary finite metrics into the line metric $\mathbb{R}$. The average distortion is the factor by which the average distance in the metric is stretched.

Related Work The definition of average distortion is not new; e.g., Alon et al. [3] study the question of embedding a metric into a tree with low average distortion. In recent work on average distortion that is closer to ours, Rabinovich [48] proves bounds on average distortion of *non-expanding* embeddings into a line and shows the close connection between this and the max-flow min-cut ratio for concurrent multicommodity flow with applications to finding quotient cuts in graphs [43].

While our problem appears similar to that of finding the *Minimum Linear Arrangement (MLA)*, for which Rao and Richa [50] gave an $O(\log n)$ approximation using the notion of spreading metrics, it is subtly different: the MLA problem involves minimizing the average stretch of the edges $\sum_{\{u,v\} \in E} |\pi(u) - \pi(v)|$ under all maps $\pi : V \to [n]$, whereas the mappings in our problem are $f : V \to \mathbb{R}$, and must ensure that $|f(u) - f(v)| \geq d(u, v)\ \forall \{u, v\} \in V \times V$.

The problem of finding *Minimum Latency tours* (a.k.a. the Traveling Repairman problem) is relevant to our discussion in terms of techniques used. In this problem, one is given a metric space (V, d) and a root depot $r \in V$; a repairman starting at r has to visit all $|V| = n$ customers, one at each node of the metric. The goal is to minimize the *average waiting time* of the customers, where the waiting time (or *latency*) of a customer is the sum of the distances of all edges traversed by the repairman before visiting this customer. There are extensions of this problem to the k-repairman case, where k repairmen start off at r, and the latency of a customer is now the time at which any one of the repairman visits this customer. The version with only one repairman is known to be NP-hard even on a tree [52], and is MAX-SNP hard in general [15]. The first constant-factor approximation for this problem was given by Blum et al. [15]; the approximation factor was improved by Goemans and Kleinberg [35] to 7.18, and most recently by Chaudhuri et al. [20] to 3.59. For the special cases of the latency problem on trees, Arora and Karakostas [7] gave a quasi-polynomial time approximation scheme (QPTAS); similar results were given for the case when the points lie in $\mathbb{R}^d$ for fixed dimension d. The k-repairmen version of the problem was studied by [29] who show a 16.994-approximation for arbitrary k; this was improved to 8.49 by [20].

Finally, a problem that combines both the cost of a tour as well as its latency into one objective function is that of finding *time dependent TSP tours*; the paper by Blum et al. [15] gives a constant factor approximation algorithm for this problem.

Our results We prove that finding the best embedding of even a tree metric into a line metric so as to minimize the average distortion is NP-hard, and hence focus on *approximating* the average distortion of the best possible embedding for the given input metric. We give a constant-factor approximation for the problem of embedding general metrics into the line metric. For the case of n-point tree metrics, we provide a quasi-polynomial time approximation scheme (QPTAS) which outputs an embedding with distortion at most $(1+\epsilon)$ times the optimum in time $n^{O(\log n/\epsilon^2)}$. We also consider the average distortion, where the average is taken only over the endpoints of the edges of an input tree metric, we show how to exploit the structure of tree metrics to give an exact solution in polynomial time.

The first idea is to think of an embedding into the line as a tour on the nodes of the original metric that starts from the leftmost vertex on the line and visit the vertices in order from left to right.

Our results build on this simple observation, and demonstrate a close relationship between minimizing average distortion and the related problems of finding short TSP tours [42], minimum latency tours [15, 35, 4], and optimal k-repairmen solutions [29]. In particular, we prove the following results for the average distortion. These results appeared in [25].

1. **Hardness for average distortion:** We prove that the problem of finding the minimum average distortion non-contracting embedding of finite metrics into the line is NP-hard, even when the input metric is a tree metric. The proof proceeds via a reduction from the Minimum Latency Problem on trees [52].

2. **Constant-factor approximations:** We give an algorithm that embeds any metric (V, d) into the line with average distortion that is within a constant of the minimum possible over all non-contracting embeddings. In fact, we prove a slightly more general bound on non-contracting embeddings into k-spiders (i.e., homeomorphs of stars with k leaves). This result uses a lower bound on the minimum average distortion of a non-contracting embedding into a k-spider in terms of the minimum k-repairmen tour [29] on the metric.

3. **QPTAS on trees:** For tree metrics on n nodes, we give an algorithm for finding a $(1+\epsilon)$-approximation to the minimum average distortion non-contracting embedding into a line in $n^{O(\log n/\epsilon^2)}$ time. Our algorithm uses a lower bound on the minimum average distortion related to the TSP tour length and latencies of appropriately chosen segments of an optimal tour. In

this way, it extends the ideas of Arora et al. [7] for minimizing latency on trees to the more general time-dependent TSPs [15], and provides a QPTAS for the latter problem as well.

4. **Poly-time algorithm for tree-edge distortion** For a tree metric as input, if the minimum average distortion is measured only over the endpoints of the edges of the tree (we call this objective the average tree-edge distortion), then we show that an embedding following a certain Euler tour of the tree is optimal. We show how to find this tour in polynomial time by dynamic programming.

Additive distortion

Next, we consider he additive distortion of embeddings into the line metric. The additive distortion is the sum of differences in all pairwise distances between the embedded and input distances. The L_p norm of additive distortion is defined as:

$$(\sum_{x,y} |\delta(x,y) - d(x,y)|^p)^{1/p}.$$

It's important to note here that we do not restrict the embedding to be non-contracting. Instead we consider the absolute difference between the distances.

Related Work The additive distortion as a measure of the quality of the embedding has received much attention, especially for the numerical taxonomy problem. The numerical taxonomy problem is one of finding a tree metric that closely fits the input metric data. Formulation of this problem as the minimization of additive distortion was first proposed by [19] in 1967. In 1977, Waterman et al. [54] showed that if there is a tree metric T coinciding exactly with the input data D, then it can be constructed in linear time. In the case when there is no tree that fits the data perfectly, Agarwala et al. [2] used the framework of approximation algorithms to give heuristics with provable guarantees for the problem. They gave a 3-approximation to the L_∞ norm of the additive distortion for fitting the data to a tree metric. They reduced the problem to that of fitting the data to *ultrametric*, where each leaf is at the same distance from a common root. For *ultrametrics*, they used an exact polynomial-time algorithm for the L_∞ norm due to Farach et al. [32].

In our setting the host metric is the line metric. The problem of finding the embedding into the line metric with minimum additive distortion was shown to be NP-hard by [51]. The special case of the problem for the L_∞ norm (i.e. with $p = \infty$) was considered by Håstad et al. [37]. They gave a 2-approximation for it.

For fitting points to a line, a well-known result due to Menger (see e.g. [21]) gives the following four point criterion. The four point criterion says that, if every subset of size 4 can be mapped into the real line exactly, then all the points can be mapped into the line exactly. An approximate version of Menger's result was given by Bădoiu et al. [11]. They proved that if every subset of size

4 can be embedded into the line with the L_∞ norm of the additive distortion being at most ϵ then all the points can be embedded with the L_∞ norm of the additive distortion being at most 6ϵ.

Our results Our main result is $O(\log^{1/2p}(n))$-approximation algorithm for the L_p norm of the additive distortion. The result is obtained in two steps. First we show that an r-fixed embedding is within a factor 3 of the optimal embedding. Here, r-fixed embedding means that there is a point r such that the distance of any other point from r is same as that in the input. Next, we 'guess' the point r and approximate the best r-fixed embedding. We show that the problem of finding the best r-fixed embedding is same as two-cost bipartition problem. We encode the two cost bipartition as a Min-Uncut problem and use $O(\sqrt{\log n})$-approximation for the latter due to Agarwal et al. [1]. This gives an $O(\sqrt{\log n})$-approximation to the p^{th} power of L_p norm, or equivalently, $O(\log^{1/2p}(n))$-approximation algorithm for the L_p norm of the additive distortion. This result has appeared in [22]

Classical Distortion

Next we address the problem of approximating the classical distortion. Given a graph $G = (V, E)$ inducing a shortest path metric $M = M(G) = (V, d)$, find a mapping f of V into a *line* that is non-contracting (i.e., $|f(u) - f(v)| \geq d(u, v)$ for all $u, v \in V$) which minimizes the distortion $D(M, f) = \max_{u,v \in V} \frac{|f(u)-f(v)|}{d(u,v)}$. That is, our goal is to find $D(M) = \min_f D(M, f)$.

Related Work Majority of the results for distortion prove uniform bounds on the distortion of embeddings. Algorithmic results on finding the best distortion are relatively few. Recently, Kenyon, Rabani and Sinclair [39] gave *exact* algorithms for minimum distortion embeddings of metrics *onto* simpler metrics (e.g., line metrics). Their algorithms work as long as the minimum distortion is small, e.g., constant. We note that constraining the embeddings to be *onto* (not *into*, as in our case) is crucial for the correctness of their algorithms.

Very recently Bădoiu et al. [9] gave an $o(n)$-approximation algorithm for embedding weighted graphs into the line metric. They also showed that it is NP-hard to approximate it within n^δ for some small constant $\delta > 0$.

Our results For the case when G is an *unweighted* graph, we show the following algorithms for this problem (denote $n = |V|$):

- A polynomial $O(D)$-approximation algorithm for metrics (V, d) for which the optimal distortion is D. This also implies an $O(\sqrt{n})$-approximation algorithm.
- A polynomial-time $\tilde{O}(\sqrt{D})$ approximation algorithm for metrics generated by unweighted trees. This also implies an $\tilde{O}(n^{1/3})$-approximation algorithm for these metrics.

These results have appeared as part of [10].

Techniques We prove two simple lower bounds on the distortion. First one is *local density*, the number of points in a set of radius r divided by r. Second one is 3-spider bound. A 3-spider is a tree where all vertices have degree at most 2, except one which has degree 3. In other words, it looks like 3 paths starting from a common point. The length of a 3-spider is the length of the shortest of the 3 paths. The 3-spider lower bound says that if the graph contains a 3-spider of length l, then the distortion is at least $\Omega(l)$.

In our algorithm, we first 'guess' the value of optimal distortion D. We divide the graph along a diametric path to get components with diameter $O(D)$. Due to the local density lower bound, the number of vertices in each component is $O(D^2)$. We find a 'good' embedding (with distortion $O(D^2)$) for each of the components and concatenate the embeddings together to get the embedding for the whole graph. This gives the total distortion to be $O(D^2)$.

In case of unweighted tree metrics, we can employ a sophisticated embedding for each component. The distortion is improved to $\tilde{O}(D^{1.5})$.

For a special case in unweighted trees, where all the subtrees in a component have size smaller than $O(D)$, we give an improved $O(\log n)$-approximation using a randomized algorithm.

Weighted Bandwidth

Finally, we consider the problem of finding a linear ordering that minimizes the stretch. In other words, given a metric (V, d), we want to map the points to $\{1, 2, \ldots, n\}$, so as to minimize $max_{x,y}|f(x) - f(y)|/d(x, y)$. i.e., instead of non-contracting embedding, we look for just a linear ordering.

Our results We give an $O(\log^2(n) \log \Delta)$-approximation for this problem. As a generalization of this result, we also get an approximation algorithm for the weighted bandwidth problem. Weighted bandwidth is defined $max_{x,y}|f(x)-f(y)|w(x, y)$, where $w(x, y)$ denotes the weight of the edge (x, y). Our approximation guarantee for the weighted bandwidth problem is $O(\log^2 n \log nW)$, where W is the maximum weight of an edge.

Techniques First we show how to reduce the weighted bandwidth problem to the problem of minimizing stretch of a linear embedding. From the weights of the edges, we construct a metric as follows: the length of an edge $l(x, y) := 1/w(x, y)$ and the distance between two points is computed using shortest path metric. We show that the stretch of this metric is same as the weighted bandwidth of original graph.

For approximating the stretch of the metric, we introduce a lower bound based on local density. The rest of our algorithm is similar to Feige's algorithm [33] for (unweighted) bandwidth problem.

These results appear in [23].

Related Work The (unweighted) bandwidth minimization problem (i.e. when all the edge weights are 1) arises in VLSI layout problems and has received much attention. It was shown to be NP-hard by Papadimitriou [47]. Blum et al. [16] gave an SDP relaxation of the bandwidth and obtained an $O(\sqrt{n/b^*})$ approximation. The first non-trivial approximation to this problem was given by Feige [33]. He developed a notion of *volume-respecting* embedding and used it to give $O(\log^{4.5} n)$-approximation for the bandwidth problem. Subsequently, Dunagan and Vempala [26] showed how to improve the approximation factor based on the SDP relaxation of Blum et al. [16]. Recently, Krauthgamer et al. [41] showed an algorithm to construct volume respecting embeddings and thus reduced the approximation factor to $O(\log^3 n)$.

Embeddings into spanning trees

In this chapter, we study probabilistic embeddings of graphs into induced spanning trees. Given a graph $G = (V, E)$, we consider the shortest path metric on it defined by (V, d). A probabilistic embedding into induced spanning trees is a probability distribution over the spanning trees of graph G. The quality of the embedding is measured by *expected distortion*.

Note that we are interested in uniform bounds on the expected distortion.

Related Work The problem of embedding a graph into a spanning tree to minimize the average distortion was first considered by Alon, Karp, Peleg and, West [3]. They gave an algorithm to construct a spanning tree with $O(exp(\sqrt{\log n \log\log n}))$ average distortion and applied to the online K-server problem. They also demonstrated examples where $\Omega(\log n)$ average distortion would be incurred for any spanning tree.

Subsequently, Bartal [12, 13] considered the problem of probabilistic embeddings of arbitrary metrics into tree metrics (not necessarily spanning trees). He obtained $O(\log^2 n)$ expected distortion and subsequently improved it to $O(\log n \log\log n)$. He also proved a lower bound of $\Omega(\log n)$ for expander graphs. Later Fakcharoenphol, Rao and Talwar gave an algorithm with $O(\log n)$ expected distortion, thus matching the lower bound.

The case of spanning trees was still open. Recently, Emek and Peleg [28] gave an $O(\log n)$-approximation algorithm for minimizing the distortion of a single spanning tree.

In 2004, Spielman and Teng [53] showed that embedding into a spanning tree with average stretch ϕ yields an $O(m\phi \log^{O(1)} n)$-time algorithm for solving diagonally-dominant symmetric linear systems. Subsequently, Elkin, Emek, Spielman and Teng [27] made a breakthrough for the average distortion problem. Their algorithm had $O(\log^2 n \log\log n)$ average distortion.

Our results We give a simple algorithm with $O(\log^2 n)$ expected distortion. This also implies an $O(\log^2 n)$ bound on average distortion. Our algorithm using the star-decomposition schema introduced by Elkin et al. [27]. We combine it with the cutting scheme of Bartal [12]. Furthermore, we introduce a new technique, viz. tail-bounds on the diameter of the resulting sub-trees to bound the distortion. Our techniques are orthogonal to those used by Elkin et al.. It might be possible to improve upon our results by combining these ideas.

These results also appear as part of [24].

References

[1] Amit Agarwal, Moses Charikar, Konstantin Makarychev, and Yury Makarychev. $o(\sqrt{\log n})$ approximation algorithms for min uncut, min 2cnf deletion, and directed cut problems. In *Proceedings of the 37th ACM Symposium on the Theory of Computing (STOC)*, pages 573–581, 2005.

[2] Richa Agarwala, Vineet Bafna, Martin Farach, Mike Paterson, and Mikkel Thorup. On the approximability of numerical taxonomy (fitting distances by tree metrics). *SIAM J. Comput.*, 28(3):1073–1085 (electronic), 1999.

[3] Noga Alon, Richard M. Karp, David Peleg, and Douglas West. A graph-theoretic game and its application to the k-server problem. *SIAM J. Comput.*, 24(1):78–100, 1995.

[4] Aaron Archer and David P. Williamson. Faster approximation algorithms for the minimum latency problem. In *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 88–96, 2003.

[5] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.

[6] Sanjeev Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. In *37th Annual Symposium on Foundations of Computer Science (Burlington, VT, 1996)*, pages 2–11. IEEE Comput. Soc. Press, Los Alamitos, CA, 1996.

[7] Sanjeev Arora and George Karakostas. Approximation schemes for minimum latency problems. *SIAM J. Comput.*, 32(5):1317–1337 (electronic), 2003.

[8] Yonatan Aumann and Yuval Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.*, 27(1):291–301, 1998.

[9] Mihai Bădoiu, Julia Chuzhoy, Piotr Indyk, and Anastasios Sidiropoulos. Low-distortion embeddings of general metrics into the line. In *Proceedings of the 37th ACM Symposium on the Theory of Computing (STOC)*, 2005.

[10] Mihai Bădoiu, Kedar Dhamdhere, Anupam Gupta, Yuri Rabinovich, Harald Räcke, R. Ravi, and Anastasios Sidiropoulos. Approximation algorithms for embeddings into low-dimensional spaces. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 119–128, 2005.

[11] Mihai Bădoiu, Piotr Indyk, and Yuri Rabinovich. Approximate algorithms for embedding metrics into low-dimensional spaces. In *Unpublished manuscript*, 2003.

[12] Yair Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *Proceedings of the 37th Symposium on the Foundations of Computer Science (FOCS)*, pages 184–193, 1996.

[13] Yair Bartal. On approximating arbitrary metrics by tree metrics. In *Proceedings of the 30th ACM Symposium on the Theory of Computing (STOC)*, pages 161–168, 1998.

[14] Yair Bartal, Avrim Blum, Carl Burch, and Andrew Tomkins. A polylog(n)-competitive algorithm for metrical task systems. In *Proceedings of the 29th ACM Symposium on the Theory of Computing (STOC)*, pages 711–719, 1997.

[15] Avrim Blum, Prasad Chalasani, Don Coppersmith, Bill Pulleyblank, Prabhakar Raghavan, and Madhu Sudan. The minimum latency problem. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 163–171. ACM Press, 1994.

[16] Avrim Blum, Goran Konjevod, R. Ravi, and Santosh Vempala. Semi-definite relaxations for minimum bandwidth and other vertex-ordering problems. *Theoret. Comput. Sci.*, 235(1):25–42, 2000. (Preliminary version in *30th STOC*, pages 100–105, 1998).

[17] Jean Bourgain. On Lipschitz embeddings of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985.

[18] Gruia Călinescu, Howard Karloff, and Yuval Rabani. Approximation algorithms for the 0-extension problem. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 8–16. ACM Press, 2001.

[19] L. Cavalli-Sforza and A. Edwards. Phylogenetic analysis models and estimation procedures. *American Journal of Human Genetics*, 19:233–257, 1967.

[20] Kamalika Chaudhuri, Brighten Godfrey, Satish Rao, and Kunal Talwar. Paths, trees and minimizing latency. In *Proceedings of the 44th Symposium on the Foundations of Computer Science (FOCS)*, pages 36–45, 2003.

[21] Michel Marie Deza and Monique Laurent. *Geometry of cuts and metrics*, volume 15 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1997.

[22] Kedar Dhamdhere. Approximating additive distortion of embeddings into line metrics. In *Proceedings of the International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, 2004.

[23] Kedar Dhamdhere. A note on minimizing weighted bandwidth of a graph. in submission, 2006.

[24] Kedar Dhamdhere, Anupam Gupta, and Harald Räcke. Improved embeddings into spanning trees. in submission, 2006.

[25] Kedar Dhamdhere, Anupam Gupta, and R. Ravi. Approximation algorithms for minimizing average distortion. In *21st Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2996 of *Lecture Notes in Computer Science*, pages 234–245, 2004.

[26] John Dunagan and Santosh Vempala. On Euclidean embeddings and bandwidth minimization. In *Approximation, randomization, and combinatorial optimization (Berkeley, CA, 2001)*, volume 2129 of *Lecture Notes in Comput. Sci.*, pages 229–240. Springer, Berlin, 2001.

[27] Michael Elkin, Yuval Emek, Dan Spielman, and Shang-Hua Teng. Lower-stretch spanning trees. In *Proceedings of the thirty-seventh ACM symposium on Theory of computing (STOC)*, 2005.

[28] Yuval Emek and David Peleg. Approximating minimum max-stretch spanning trees on unweighted graphs. In *SODA '04: Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 261–270, Philadelphia, PA, USA, 2004. Society for Industrial and Applied Mathematics.

[29] Jittat Fakcharoenphol, Chris Harrelson, and Satish Rao. The k-traveling repairman problem. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 655–664. Society for Industrial and Applied Mathematics, 2003.

[30] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the thirty-fifth ACM symposium on Theory of computing*, pages 448–455. ACM Press, 2003.

[31] M. Farach, S. Kannan, and T. Warnow. A robust model for finding optimal evolutionary trees. *Algorithmica*, 13(1-2):155–179, 1995.

[32] M. Farach, S. Kannan, and T. Warnow. A robust model for finding optimal evolutionary trees. *Algorithmica*, 13:155–179, 1995.

[33] Uriel Feige. Approximating the bandwidth via volume respecting embeddings. *J. Comput. System Sci.*, 60(3):510–539, 2000. 30th Annual ACM Symposium on Theory of Computing (Dallas, TX, 1998).

[34] Naveen Garg, Goran Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group Steiner tree problem. *Journal of Algorithms*, 37(1):66–84, 2000. (Preliminary version in *9th SODA*, pages 253–259, 1998).

[35] Michel Goemans and Jon Kleinberg. An improved approximation ratio for the minimum latency problem. *Math. Programming*, 82(1-2, Ser. B):111–124, 1998. Networks and matroids; Sequencing and scheduling.

[36] Anupam Gupta. Steiner points in tree metrics don't (really) help. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (Washington, DC, 2001)*, pages 220–227, Philadelphia, PA, 2001. SIAM.

[37] Johan Håstad, Lars Ivansson, and Jens Lagergren. Fitting points on the real line and its application to RH mapping. In *Algorithms—ESA '98 (Venice)*, volume 1461 of *Lecture Notes in Comput. Sci.*, pages 465–476. Springer, Berlin, 1998.

[38] Piotr Indyk. Algorithmic aspects of geometric embeddings. In *Proceedings of the 42nd Symposium on the Foundations of Computer Science (FOCS)*, pages 10–33, 2001.

[39] Claire Kenyon, Yuval Rabani, and Alistair Sinclair. Low distortion maps between point sets. In *Proceedings of the 36th ACM Symposium on the Theory of Computing (STOC)*, pages 272–380, 2004.

[40] Jon Kleinberg and Éva Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields. *Journal of the ACM (JACM)*, 49(5):616–639, 2002. (Preliminary version in *40th FOCS*, 1999).

[41] Robert Krauthgamer, James Lee, Manor Mendel, and Assaf Naor. Measured descent: A new embedding method for finite metrics. In *Proceedings of the 45th Symposium on the Foundations of Computer Science (FOCS)*, pages 434–443, 2004.

[42] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and editors D. B. Shmoys. *The Traveling Salesman Problem*. John Wiley & Sons, 1985.

[43] F. Thomson Leighton and Satish B. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, 1999. (Preliminary version in *29th FOCS*, pages 422–431, 1988).

[44] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. (Preliminary version in *35th FOCS*, 1994).

[45] Jiří Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.

[46] Jiří Matoušek. Bi-Lipschitz embeddings into low dimensional Euclidean spaces. *Commentationes Mathematicae Universitatis Carolinae*, 31(3):589–600, 1990.

[47] Christos H. Papadimitriou. The NP-completeness of the bandwidth minimization problem. *Computing*, 16(3):263–270, 1976.

[48] Yuri Rabinovich. On average distortion of embedding metrics into the line and into ℓ_1. In *Proceedings of the thirty-fifth ACM symposium on Theory of computing*, pages 456–462. ACM Press, 2003.

[49] Yuri Rabinovich and Ran Raz. Lower bounds on the distortion of embedding finite metric spaces in graphs. *Discrete Comput. Geom.*, 19(1):79–94, 1998.

[50] Satish Rao and Andrea Richa. New approximation techniques for some ordering problems. In *Proceedings of the 9th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 211–218, 1998.

[51] James B. Saxe. Embeddability of graphs into k-space is strongly np-hard. In *Allerton Conference in Communication, Control and Computing*, pages 480–489, 1979.

[52] R.A. Sitters. The minimum latency problem is np-hard for weighted trees. In *W.J. Cook, A.S. Schulz (eds.), Integer Programming and Combinatorial Optimization, Lecture Notes in Computer Science 2337*, pages 230–239, 2002.

[53] Dan Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th ACM Symposium on the Theory of Computing (STOC)*, 2004.

[54] M. S. Waterman, T. S. Smith, M. Singh, and W. A. Beyer. Additive evolutionary trees. *Journal of Theoretical Biology*, 64:199–213, 1977.