

Low-Distortion Embeddings of Trees

Robert Babilon Jiří Matoušek Jana Maxová Pavel Valtr

Department of Applied Mathematics and
Institute for Theoretical Computer Science (ITI),
Charles University, Malostranské nám. 25,
118 00 Prague, Czech Republic
{babilon,matousek,jana,valtr}@kam.mff.cuni.cz

Abstract

We prove that every tree $T = (V, E)$ on n vertices with edges of unit length can be embedded in the plane with distortion $O(\sqrt{n})$; that is, we construct a mapping $f: V \rightarrow \mathbf{R}^2$ such that $\rho(u, v) \leq \|f(u) - f(v)\| \leq O(\sqrt{n}) \cdot \rho(u, v)$ for every $u, v \in V$, where $\rho(u, v)$ denotes the length of the path from u to v in T . The embedding is described by a simple and easily computable formula. This is asymptotically optimal in the worst case. We also construct interesting optimal embeddings for a special class of trees (fans consisting of paths of the same length glued together at a common vertex).

Communicated by: P. Mutzel and M. Jünger;
submitted May 2002; revised April 2003.

1 Introduction

Embeddings of finite metric spaces into Euclidean spaces or other normed spaces that approximately preserve the metric received considerable attention in recent years. Numerous significant results have been obtained in the 1980s, mainly in connection with the local theory of Banach spaces. Later on, surprising algorithmic applications and further theoretical results were found in theoretical computer science; see, for example, Indyk [2] or Chapter 15 in [4] for overviews.

Here we consider only embeddings into the spaces \mathbf{R}^d with the Euclidean metric. The quality of an embedding is measured by the *distortion*. Let (V, ρ) be a finite metric space. We say that a mapping $f: V \rightarrow \mathbf{R}^d$ has distortion at most D if there exists a real number $\alpha > 0$ such that for all $u, v \in V$,

$$\alpha \cdot \rho(u, v) \leq \|f(u) - f(v)\| \leq D \cdot \alpha \cdot \rho(u, v).$$

This definition permits scaling of all distances in the same ratio α , in addition to the distortion of the individual distances by factors between 1 and D . Since the image in \mathbf{R}^d can always be re-scaled as needed, we can choose the factor α at our convenience.

We study embeddings with the dimension d fixed, and mainly the case $d = 2$, i.e. embeddings into the plane. A straightforward volume argument shows that the distortion required to embed into \mathbf{R}^d the n -point metric space with all distances equal to 1 is at least $\Omega(n^{1/d})$. In [3], it was proved that general metric spaces may require even significantly bigger distortions; namely, for every fixed d , there are n -point metric spaces that need $\Omega(n^{1/\lfloor (d+1)/2 \rfloor})$ distortion for embedding into \mathbf{R}^d .


Better upper bounds can be obtained for special classes of metric spaces. A metric ρ on a finite set V is called a *tree metric* if there is a tree $T = (V', E)$ (in the graph-theoretic sense) with $V' \supseteq V$ and a weight function $w: E \rightarrow (0, \infty)$ such that for all $u, v \in V$, $\rho(u, v)$ equals the length of the path connecting the vertices u and v in T , where the length of an edge $e \in E$ is $w(e)$. It was conjectured in [3] that, for fixed d , all n -point tree metrics can be embedded into \mathbf{R}^d with distortion $O(n^{1/d})$. If true, this is best possible, as the example of a star with $n - 1$ leaves and unit-length edges shows (a volume argument applies). Gupta [1] proved the somewhat weaker upper bound $O(n^{1/(d-1)})$.

Here we make a step towards establishing the conjecture. We deal with the planar case, where the gap between the lower bound of \sqrt{n} and Gupta's $O(n)$ upper bound is the largest (in fact, the $O(n)$ upper bound holds even for general metric spaces and embeddings into \mathbf{R}^1 [3]). So far we can only handle the case of unit-length edges.

Theorem 1 *Every n -vertex tree with unit-length edges, considered as a metric space, can be embedded into \mathbf{R}^2 with distortion $O(\sqrt{n})$. The embedding is described by a simple explicit formula and can be computed efficiently.*

Gupta's result actually states that if the considered tree has at most ℓ leaves, then an embedding into \mathbf{R}^d with distortion $O(\ell^{1/(d-1)})$ is possible. We show

that here the dependence on ℓ *cannot* be improved. Let $F_{\ell,m}$ (the fan with ℓ leaves and path length m) denote the tree consisting of ℓ paths of length m each glued together at a common vertex (root):



Proposition 1 *For every fixed $d \geq 2$ and every $m, \ell \geq 2$ every embedding of $F_{\ell,m}$ into \mathbf{R}^d requires distortion $\Omega(\ell^{1/(d-1)})$ if $\ell \leq m^{\frac{d(d-1)}{d+1}}$, and distortion $\Omega(\ell^{1/d}m^{1/(d+1)})$ if $\ell \geq m^{\frac{d(d-1)}{d+1}}$.*

Theorem 1 provides embeddings which are optimal in the worst case, i.e. there are trees for which the distortion cannot be asymptotically improved. But optimal or near-optimal embeddings of special trees seem to present interesting challenges, and sometimes low-distortion embeddings are aesthetically pleasing and offer a good way of drawing the particular trees. We present one interesting example concerning the fan $F_{\sqrt{n},\sqrt{n}}$. Here the “obvious” embedding as in the above picture, as well as the embedding from Theorem 1, yield distortions $O(\sqrt{n})$. In Section 4, we describe a somewhat surprising better embedding, with distortion only $O(n^{5/12})$. This is already optimal according to Proposition 1.

Several interesting problems remain open. The obvious ones are to extend Theorem 1 to higher dimensions and/or to trees with weighted edges. Another, perhaps more difficult, question is to extend the class of the considered metric spaces. Most significantly, we do not know an example of an n -vertex planar graph (with weighted edges) whose embedding into \mathbf{R}^2 would require distortion larger than about \sqrt{n} . If it were possible to show that all planar-graph metrics can be embedded into the plane with $o(n)$ distortion, it would be a neat metric condition separating planar graphs from non-planar ones, since suitable n -vertex subdivisions of any fixed non-planar graph require $\Omega(n)$ distortion [3].

2 Proof of Theorem 1

Notation.

Let T be a tree (the edges have unit lengths) and let ρ be the shortest-path metric on $V = V(T)$. One vertex is chosen as a root. The *height* $h(v)$ of a vertex $v \in V$ is its distance to the root. Let π_v denote the path from a vertex v to the root.

For every vertex we fix a linear (left-to-right) ordering of its children. This defines a partial ordering \preceq on V : we have $u \prec v$ iff $u \notin \pi_v$, $v \notin \pi_u$, and π_u

goes right of π_u at the vertex where π_u and π_v branch. We define

$$\text{sgn}_v(u) = \begin{cases} 0 & \text{if } u \in \pi_v \text{ or } v \in \pi_u \\ +1 & \text{if } u \prec v \\ -1 & \text{if } v \prec u. \end{cases}$$

Furthermore we define $a_v(u)$ as the distance of v to the nearest common ancestor of u and v , and $\ell(v) = |\{u \in V(T) : u \prec v\}|$.

Construction.


Our construction resembles Gupta’s construction [1] to some extent, but we needed several new ideas to obtain a significantly better distortion.

First we describe the idea of the construction.

The embedded tree is placed in a suitable acute angle. The apex of the angle is at the origin of coordinates and the angle opens upwards. The root of the tree is placed to the origin.

In the first step we put vertices of height h on the horizontal line $y = 2\sqrt{nh}$. Our angle is divided into subangles that correspond to subtrees rooted at the sons of the root; the sizes of these subangles are proportional to the number of vertices in the subtrees. The vertices of height 1 are placed on the first horizontal line in the middle of the corresponding subangles. We continue in the same way for every subtree, with each subangle translated so that its vertex lies at the root of the subtree.

In the second (and last) step we raise the y -coordinates of the vertices by appropriate amounts between 0 and \sqrt{n} .



Now we describe the embedding more formally; actually we use a slightly modified embedding for which the formulas are quite simple.

First we define an auxiliary function $k(v)$:

$$k(v) = \sum_{w \in V} \operatorname{sgn}_v(w) a_v(w).$$

The embedding $f: V \rightarrow \mathbf{R}^2$ is now defined by $f(v) = (x(v), y(v))$, where

$$\begin{aligned} x(v) &= n^{-1/2} k(v) \\ y(v) &= 2\sqrt{n} h(v) + (\ell(v) \bmod \sqrt{n}). \end{aligned}$$

Distortion.

We now estimate the distortion of the embedding defined above. First we show that the maximum expansion of the tree distance of any two vertices is at most $O(\sqrt{n})$. By the triangle inequality, the maximal expansion is always attained on an edge. Hence, it suffices to consider two vertices connected by an edge. So let u be the father of v in T ; it suffices to verify that $\|f(u) - f(v)\| = O(\sqrt{n})$. We clearly have $|y(u) - y(v)| \leq 3\sqrt{n}$, and so it suffices to prove that $|x(u) - x(v)| = O(\sqrt{n})$, which will follow from $|k(u) - k(v)| \leq n$. We have

$$|k(v) - k(u)| \leq \sum_{w \in V(T)} \left| \operatorname{sgn}_u(w) a_u(w) - \operatorname{sgn}_v(w) a_v(w) \right|.$$

We claim that every w contributes at most 1 to this sum. First we note that $|a_u(w) - a_v(w)| \leq 1$. If $\operatorname{sgn}_u(w) \neq 0$, we have $\operatorname{sgn}_v(w) = \operatorname{sgn}_u(w)$, and the contribution of w is at most $|a_u(w) - a_v(w)| \leq 1$. The only remaining possibility is $\operatorname{sgn}_u(w) = 0$ and $\operatorname{sgn}_v(w) \neq 0$. But then u is the nearest common ancestor of v and w , we have $a_v(w) = 1$, $a_u(w) = 0$, and the contribution of w is at most 1 in this case as well. Therefore $|k(u) - k(v)| \leq n$ as claimed.

Next, we are going to prove that $\|f(u) - f(v)\| = \Omega(\rho(u, v))$ for every $u, v \in V$; this will finish the proof of Theorem 1.

First we consider the situation when $h(u) = h(v)$; this is the main part of the proof, and the case where u and v have different levels will be an easy extension of this. So let $h(u) = h(v)$ and $a = a_v(u) = a_u(v)$, and assume that $u \prec v$.

Lemma 1 $\ell(v) - \ell(u) \geq a$.

This is because all the vertices on π_u not lying on π_v are counted in $\ell(v)$ but not in $\ell(u)$. □

Corollary 1 *If $\ell(v) - \ell(u) < \frac{1}{2}\sqrt{n}$, then $|y(u) - y(v)| \geq \frac{1}{2}\rho(u, v)$.*

Indeed, we have $|y(u) - y(v)| = |(\ell(u) \bmod \sqrt{n}) - (\ell(v) \bmod \sqrt{n})| \geq |\ell(u) - \ell(v)| \geq a = \frac{1}{2}\rho(u, v)$. □

Thus, the difference in the y -coordinate takes care of u and v whenever $\ell(v) - \ell(u) < \frac{1}{2}\sqrt{n}$. Next, we need to show that if this is not the case, then the x -coordinate takes care of u and v .

Lemma 2 $k(v) - k(u) \geq [\ell(v) - \ell(u)] \cdot a$.

Corollary 2 *If $\ell(v) - \ell(u) \geq \frac{1}{2}\sqrt{n}$, then $x(v) - x(u) \geq n^{-1/2}[\ell(v) - \ell(u)] \cdot a \geq \frac{1}{2}a \geq \frac{1}{4}\rho(u, v)$.* \square

Combining this corollary with Corollary 1 yields $\|f(u) - f(v)\| = \Omega(\rho(u, v))$ for all $u, v \in V$ with $h(u) = h(v)$.

Proof of Lemma 2: We have $k(v) - k(u) = \sum_{w \in V} t_{u,v}(w)$, where $t_{u,v}(w) = \text{sgn}_v(w)a_v(w) - \text{sgn}_u(w)a_u(w)$. First we check that $t_{u,v} \geq 0$ for all w . We always have $\text{sgn}_v(w) \geq \text{sgn}_u(w)$, and so the only cases that might cause trouble are $\text{sgn}_v(w) = \text{sgn}_u(w) = 1$ and $\text{sgn}_v(w) = \text{sgn}_u(w) = -1$. We will analyze only the case when $\text{sgn}_v(w) = \text{sgn}_u(w) = 1$, the other one is symmetric. In this case $w \prec u \prec v$, and it is easy to check that then $a_v(w) \geq a_u(w)$, which shows $t_{u,v}(w) \geq 0$ in all cases.

Next, we verify that if w contributes 1 to $\ell(v) - \ell(u)$, which means $w \prec v$ and $w \not\prec u$, then $t_{u,v}(w) \geq a$. We distinguish two cases for such w : either $u \prec w \prec v$, or w and u lie on a common path to the root, i.e. $w \in \pi_u$ or $u \in \pi_w$.

In the first case, $u \prec w \prec v$, $\text{sgn}_u(w) = -1$ and $\text{sgn}_v(w) = 1$, and so $t_{u,v}(w) = a_u(w) + a_v(w)$. It is easy to see that if u and v are vertices with $a_u(v) = a_v(u) = a$, then $\max(a_u(w), a_v(w)) \geq a$ for all $w \in V$ and so $t_{u,v}(w) \geq a$ in this case.

In the second case, we have $\text{sgn}_v(w) = 1$ and $\text{sgn}_u(w) = 0$, and so $t_{u,v}(w) = a_v(w) = a$, since the nearest common ancestor of u and v is the same as the nearest common ancestor of w and v . The proof of Lemma 2 is complete. \square

Claim 1 *If two vertices u, v are on different levels, then*

$$\|f(v) - f(u)\| \geq \frac{1}{12\sqrt{2}} \rho(u, v).$$

Proof: Without loss of generality, suppose that $h(u) < h(v)$ and $u \prec v$. Set $a := a_u(v)$, and let w be the vertex on the path from u to v with $h(w) = h(u)$, $\rho(u, w) = 2a$.

If $\rho(u, v) > 3a$, then $\|f(v) - f(u)\| \geq y(v) - y(u) \geq (\rho(u, v) - 2a)\sqrt{n} \geq \frac{\sqrt{n}}{3}\rho(u, v) \geq \frac{1}{3}\rho(u, v)$.

If $\rho(u, v) \leq 3a$ and $a < 4\sqrt{n}$, then $\|f(v) - f(u)\| \geq y(v) - y(u) \geq \sqrt{n} > a/4 \geq \frac{1}{12}\rho(u, v)$.

Finally, suppose that $\rho(u, v) \leq 3a$ and $a \geq 4\sqrt{n}$. Using the definition of the embedding f and with the observation that $k(v) \leq nh(v)$, it is easy to check that $|x(v)| \leq y(v)$ for any vertex v . Thus the whole drawing of the tree lies in the cone bounded by the two rays emanating from the root diagonally upwards and spanning the angle of 45° with the y -axis. This also holds for every subtree. So the distance of the point $f(u)$ to $f(v)$ is larger than the distance to the left arm of the cone rooted at $f(w)$, which is $\frac{1}{\sqrt{2}}(x(w) - x(u) + y(w) - y(u)) > \frac{1}{\sqrt{2}}(x(w) - x(u) - \sqrt{n}) \geq \frac{1}{\sqrt{2}}(a/2 - \sqrt{n}) \geq \frac{1}{\sqrt{2}} \frac{a}{4} \geq \frac{1}{12\sqrt{2}}\rho(u, v)$ (in the second inequality we used Corollary 2). Thus, $\|f(v) - f(u)\| > \frac{1}{12\sqrt{2}}\rho(u, v)$ in this case. \square




Figure 1: Illustration to the proof of Proposition 1.

3 Proof of Proposition 1

Let $F_{\ell,m}$ be defined as in Section 1; that is, $F_{\ell,m}$ denotes the tree consisting of ℓ paths of length m glued at the root r . We show that every embedding of $F_{\ell,m}$ into \mathbf{R}^d requires distortion $\Omega(\ell^{1/(d-1)})$ if $\ell \leq m^{\frac{d(d-1)}{d+1}}$, and distortion $\Omega(\ell^{1/d}m^{1/(d+1)})$ if $\ell \geq m^{\frac{d(d-1)}{d+1}}$.

Let $f: V(F_{\ell,m}) \rightarrow \mathbf{R}^d$ be any non-contracting embedding with distortion D (i.e. for all $u, v \in V$, $\rho(u, v) \leq \|f(u) - f(v)\| \leq D \cdot \rho(u, v)$). We choose a real number R such that there are $\ell/2$ of the images of endvertices of $F_{\ell,m}$ contained inside the ball B of radius R around $f(r)$ or on its boundary, and there are $\ell/2$ of them outside the ball or on its boundary.

We consider a ball of radius m around each of the $\ell/2$ images of endvertices contained in B . These balls are disjoint and contained in m -neighbourhood of B . Hence, by a volume argument, $R = \Omega(m\ell^{1/d})$.

Let S denote the sphere of radius R around $f(r)$ (i.e. the boundary of B). Let $P_1, \dots, P_{\ell/2}$ denote the $\ell/2$ paths for which the images of their endvertices are not contained inside B . Each $f(P_i)$ intersects S at least once (by $f(P_i)$ we mean $f(V(P_i))$ together with the straight lines between any $f(u), f(v)$ such that $\{u, v\} \in E(P_i)$). Let m_i denote the number of vertices on $f(P_i)$ inside B before the first intersection with S . We put $\mu = \min\{m_i \mid i = 1, \dots, \ell/2\}$. From the choice of μ it follows that $R \leq D\mu$.

Let P be a spherical shell of width D around B ; that is,

$$P = B(f(r), R + D) \setminus \text{int } B(f(r), R),$$

where $\text{int}X$ stands for the interior of X . At least $\ell/2$ vertices (one from each path

$P_i, i = 1, \dots, \ell/2$ are contained in P . We consider a ball of radius μ around each of these vertices. These balls are disjoint (because the tree-distance between any two of the vertices is at least 2μ) and they are contained in μ -neighbourhood of P (see the picture). Therefore we have, for a suitable constant $c > 0$,

$$R^{d-1}(D + 2\mu) \geq c\ell\mu^d.$$

Thus there are two possibilities: $R^{d-1}D = \Omega(\ell\mu^d)$ or $R^{d-1}\mu = \Omega(\ell\mu^d)$.

If $R^{d-1}D = \Omega(\ell \cdot \mu^d)$, then $D^{d+1} = \Omega(\ell \cdot R)$, since $R \leq D\mu$. And since $R = \Omega(m\ell^{1/d})$ we have $D^{d+1} = \Omega(\ell \cdot m \cdot \ell^{1/d})$. Thus $D = \Omega(\ell^{1/d} \cdot m^{1/(d+1)})$. If $R^{d-1}\mu = \Omega(\ell \cdot \mu^d)$, then since $R \leq D\mu$, we immediately have $D = \Omega(\ell^{1/(d-1)})$. Thus we know that $D \geq \min\{\Omega(\ell^{1/d} \cdot m^{1/(d+1)}), \Omega(\ell^{1/(d-1)})\}$. It means that $D = \Omega(\ell^{1/d} \cdot m^{1/(d+1)})$ if $\ell \geq m^{\frac{d(d-1)}{d+1}}$, and $D = \Omega(\ell^{1/(d-1)})$ if $\ell \leq m^{\frac{d(d-1)}{d+1}}$.

4 A low-distortion embedding of a subdivided star

In this section we describe an asymptotically optimal embedding of $F_{\sqrt{n}, \sqrt{n}}$ into \mathbf{R}^2 , with distortion $O(n^{5/12})$. The other $F_{\ell, m}$ can be embedded similarly, with distortion as in Proposition 1; we omit a proof since it is technical and contains no new ideas.

The embedding is sketched in Fig 2. The tree $F_{\sqrt{n}, \sqrt{n}}$ is embedded into the shaded trapezoids (there are $n^{1/4} \cdot 2n^{1/3}$ of them, each of height $n^{5/12}$) and into the shaded discs (there are $n^{1/4} \cdot n^{1/4}$ of them). The root of $F_{\sqrt{n}, \sqrt{n}}$ is embedded to the bottom vertex of the triangle.

In the following, we describe the embedding in more detail. It will be done in two steps. In the first step, we embed vertices up to the level $h = 2n^{1/3}$ into the dashed trapezoids on the corresponding level. In the second step, we embed all the other vertices on diameters of the discs and change the positions of some vertices on levels between $n^{1/3}$ and $2n^{1/3}$.

Step 1.

We start with the vertices on level $h = 2n^{1/3}$. There are $n^{1/2}$ vertices on this level. We divide them into $n^{1/4}$ groups, each of size $n^{1/4}$. Each of these groups is packed into one of the dashed trapezoids (on the last level). This is possible, since the area of each trapezoid on this level is $\Theta(n^{11/12})$, which is exactly the required area (we have $n^{1/4}$ vertices, and any two of them must be at distance at least $\Omega(n^{1/3})$). It follows that no two vertices on level $h = 2n^{1/3}$ are embedded too close to each other.

When all the vertices on the level $h = 2n^{1/3}$ are embedded, we embed the vertices on lower levels. We connect every vertex x on the last level by a straight line to the root. On this line we embed all the vertices on the path from x to the root. We do this in such a way that the distance Δ (measured on the connecting




Figure 2: The embedding of $F_{\sqrt{n}, \sqrt{n}}$.




Figure 3: The placement of vertices in lower levels.

line) from a vertex to the bottom of its dashed trapezoid is the same for all the vertices on the path from x to the root (see Fig. 3). It is easy to check that the distance between the vertices on the last level guarantees that the vertices in the lower levels are also far enough from each other. Indeed, if we consider two vertices x_1 and y_1 of level h_1 (i.e., with tree distance $2h_1$) that have distance d

in the embedding, then for any x_2 on the path from x_1 to the root and any y_2 on the path from y_1 to the root, both at a level h_2 , their distance is at least dh_2/h_1 (while their tree-distance is $2h_2$). Hence, if every two vertices in a higher level are far enough then also every two vertices in a lower level are far enough.

Thus, the construction described so far gives an embedding of $F_{\sqrt{n}, 2n^{1/3}}$ with distortion $O(n^{5/12})$. In the next step we modify this embedding into an embedding of $F_{\sqrt{n}, \sqrt{n}}$.

Step 2.

In this step we embed all vertices on the remaining levels. We also change the positions of some vertices on the levels between $n^{1/3}$ and $2n^{1/3}$. There is a column of discs on the left of each column of trapezoids; there are $n^{1/4}$ columns of discs with $n^{1/4}$ discs in each column. Each path of $F_{\sqrt{n}, \sqrt{n}}$ starts at the root, goes through some number of dashed trapezoids, having one vertex in each of them, reaches a suitable level, and then it turns into one of the discs, where the rest of it is embedded with vertices packed as tightly as possible. A more precise description follows.

In each column, we consider the discs one by one from the bottom to the top. For each disc, we consider the highest shaded trapezoid lying below the equator of the disc. We take the leftmost vertex embedded in that trapezoid, and we change the embedding of the successors of this vertex: The successors are embedded on the equator of the considered disc and on a straight line connecting the leftmost vertex to the equator of the considered disc. The distance between any two neighbours on the connecting line is $n^{5/12}$ and the distance between any two neighbours on the equator is 1 in the embedding.

We continue with the next higher disc (with a smaller number of paths). In the last level, there is only one path to end. The resulting embedding is indicated in Fig 2.

It is straightforward to check that the resulting embedding of $F_{\sqrt{n}, \sqrt{n}}$ yields distortion $O(n^{5/12})$ as claimed. \square

Acknowledgments

We would like to thank Helena Nyklová, Petra Smolíková, Ondřej Pangrác, Robert Šámal, and Tomáš Chudlarský for useful discussions on the problems considered in this paper. We also thank two anonymous referees for careful reading and useful comments.

References

- [1] A. Gupta. Embedding tree metrics into low dimensional Euclidean spaces. *Discrete Comput. Geom.*, 24:105–116, 2000.
- [2] P. Indyk. Algorithmic applications of low-distortion embeddings. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, 2001.
- [3] J. Matoušek. Bi-Lipschitz embeddings into low-dimensional Euclidean spaces. *Comment. Math. Univ. Carolinae*, 31:589–600, 1990.
- [4] J. Matoušek. *Lectures on Discrete Geometry*. Springer, New York, 2002.