

Combinatorial Theorems about Embedding Trees on the Real Line *

Amit Chakrabarti
Department of Computer Science
Dartmouth College
Hanover, NH 03755
ac@cs.dartmouth.edu

Subhash Khot
Computer Science Department
Courant Institute of Mathematical Sciences
New York University
New York, NY 10012
khot@cs.nyu.edu

Abstract

We consider the combinatorial problem of embedding the metric defined by an unweighted graph into the real line, so as to minimize the distortion of the embedding. This problem is inspired by connections to Banach space theory and to computer science.

After establishing a framework in which to study line embeddings, we focus on metrics defined by three specific families of trees: complete binary trees, fans, and combs. We construct asymptotically optimal (i.e., distortion-minimizing) line embeddings for these metrics and prove their optimality via suitable lower bound arguments. We show that even such specialized metrics require nontrivial constructions and proofs of optimality require sophisticated combinatorial arguments.

Our results about these metrics show that the local density of a graph — an *a priori* reasonable lower bound on the optimum distortion — might in fact be arbitrarily smaller than the true optimum, even for tree metrics. They also show that the optimum distortion for a general tree can be arbitrarily low or high, even when it has bounded degree. The combinatorial techniques from our work might prove useful in further algorithmic research on low distortion metric embeddings.

1 Introduction

Bi-Lipschitz embeddings between metric spaces have been studied in Banach space theory and, more recently, in computer science, thanks to a number of significant algorithmic applications. A basic goal in the study is to embed a finite metric space into a given target space, preserving all distances up to a small factor, called the *distortion* of the embedding (for a precise definition, see Section 2). Embeddings with low distortion into low-dimensional spaces ℓ_p^d (defined as \mathbb{R}^d equipped with the p -norm) are particularly important, their study having led to the best known algorithms for flow and cut problems [17, 19, 1], nearest neighbor searching [13] and clustering [19]. For more background and further examples of algorithmic applications, we refer the reader to the surveys by Indyk [12] and Matoušek [21, chap. 15].

We consider some especially basic questions on this subject. We focus on a simple subclass of finite metrics, *tree metrics*, defined as the shortest path metrics of unweighted trees. Further, we focus on embeddings of tree metrics into the most basic ℓ_p -space: the real line, \mathbb{R} . We call such embeddings *line embeddings*. As our work here will show, even this special case contains considerable complexity and leads to some interesting combinatorics.

*Work supported in part by NSF Grant CCR-96-23768, ARO Grant DAAH04-96-1-0181 and two NSF CAREER Awards.

Embeddings of graphs into a line have also been considered for the problem of minimizing *bandwidth* (also known as *dilation*); see, e.g., Feige [8] and the references therein. However, the notion of bandwidth ignores the metric defined by the graph and is thus quite different from the concept we study here. We say more on this below.

1.1 Our Results

This work is primarily concerned with the quantity $D^*(G)$, the least possible distortion of a line embedding of the metric defined by the (unweighted) graph G . For precise definitions, please see Section 2.

To set the background for our results, we first establish the following easy “folklore” theorems. We show that any connected n -vertex graph has a line embedding with distortion $O(n)$,¹ achieved by a depth first ordering of the vertices (Theorem 8); we think of this as the *naïve embedding* of a graph. This upper bound is easily shown to be asymptotically optimal, even when restricted to tree metrics, using the theorem $D^*(G) = \Omega(\text{ld}(G))$, where $\text{ld}(G)$ is the *local density* of G (Definition 6 and Theorem 7).

Our major theorems in this work explore the following question: how tight is the local density lower bound? It is not hard to show that for the n -cycle C_n we have $\text{ld}(C_n) = O(1)$, but $D^*(C_n) = \Omega(n)$, so the bound can be very weak. However, this leaves the question open for the simpler class of *tree* metrics. For a complete binary tree on n vertices, we exhibit a nontrivial line embedding with distortion $O(n/\log n)$, matching the local density bound (Theorem 15), whereas the naïve embedding only yields an $O(n)$ upper bound. However, we then show that the local density bound can be far from good in general, by considering two special families of trees called *fans* and *combs*. Let $\text{FAN}_{a,b}$ and $\text{COMB}_{a,b}$ denote, respectively, the fan and the comb with a arms, each of length b (examples illustrated in Fig. 1). These graphs are readily seen to have local density $O(a)$. We show, however, that $D^*(\text{COMB}_{a,b}) = \Omega(b)$ (Theorem 19) and that $D^*(\text{FAN}_{a,b}) = \Omega(\max\{b, a\sqrt{b}\})$ (Theorem 24). These lower bounds can clearly be arbitrarily higher than the local density bounds for appropriate settings of a and b . The proofs of these lower bounds are our most novel and technically challenging contributions.

We also exhibit embeddings to establish that the last two lower bounds are asymptotically optimal (Theorems 16 and 17). The embedding of $\text{FAN}_{a,b}$ is nontrivial: the naïve embedding yields only a trivial $O(ab)$ upper bound. This shows, for instance, that for $a = b = \sqrt{n}$, both the naïve embedding upper bound of $O(n)$ and the local density lower bound of $\Omega(\sqrt{n})$ are far from the truth: $D^*(\text{FAN}_{\sqrt{n},\sqrt{n}}) = \Theta(n^{3/4})$.

The large gap (i.e., ratio) between distortion and local density that we establish here is in sharp contrast to the situation for bandwidth. The bandwidth $\text{bw}(G)$ of an undirected graph G is defined as follows: $\text{bw}(G) := \min_{\pi: V(G) \leftrightarrow \{1, \dots, n\}} \max_{\{u,v\} \in E(G)} |\pi(u) - \pi(v)|$. (Note that π is a bijection.) A longstanding conjecture (see, e.g., [18, Open Problem 2]) says that $\text{bw}(G) = O(\text{ld}(G) \log n)$, where $n = |V(G)|$. Feige [8] showed, in a landmark paper, that this gap is at most $O((\log^{3.5} n) \sqrt{\log \log n})$; subsequently, Krauthgamer et al. [15] improved the bound to $O(\log^{3.5} n)$. Gupta [10] showed that the gap is at most $O(\log^{2.5} n)$ when G is either a tree or a chordal graph.

1.2 Relation to Previous Work

Within the field of metric embeddings there has been plenty of work focusing specifically on tree metrics. It is known that tree metrics admit embeddings of much better quality than general finite metric spaces. For example, an n -point tree metric embeds isometrically into ℓ_1 and into $\ell_\infty^{O(\log n)}$ [19]. In contrast, for general n -point metrics, the lowest distortion we can guarantee for embedding into ℓ_1 is only $O(\log n)$, via

¹This $O(n)$ upper bound also applies to *weighted* graphs, as shown by Matousek [20].

Bourgain’s embedding [3], and the lowest dimension we can guarantee for an isometric embedding into ℓ_∞ is only $O(n)$ [19]. Moreover, an n -point tree metric can be embedded into ℓ_2^d with distortion $\tilde{O}(n^{1/(d-1)})$ [9], whereas for general n -point metrics we can only guarantee a $O(n^{2/d})$ distortion [20].

A focus on certain specific families of trees, as in this work, is also not new to the field. For instance, Bourgain [4] considered the problem of embedding complete binary trees into ℓ_p . Diks [7] and Heckmann et al. [11] focused on binary trees and complete binary trees respectively in their work on the bandwidth problem. Bern et al. [2] considered line embeddings of complete binary trees, cycles and stars in their work on minimizing total embedded edge length.

Another related line of work is on efficient *algorithms* to approximate $D^*(G)$ for a given input graph G . Determining $D^*(G)$ exactly is an NP-hard problem [14]. Although nontrivial polynomial time approximation algorithms are known, the factors they achieve are quite large: the current best results give an $O(n^{1/2})$ -approximation for general unweighted graph metrics [5] and an $O(n^{1/3})$ -approximation for tree metrics [6]. It is also known that there is a constant $A > 1$ such that $D^*(G)$ is NP-hard to A -approximate [6]. No such hardness of approximation result is known for tree metrics. However, for *weighted* tree metrics, it is known that approximating to within $O(n^{1/12})$ is NP-hard [5].

Our work here (especially, the lower bounds) gives some additional explanation for why $D^*(G)$ is so hard to approximate well, even for tree metrics: there is more to it, combinatorially, than meets the eye, even for “simple-looking” trees. The combinatorial insights from our work might prove useful in future algorithmic research on the problem. They may also generate ideas for narrowing the large gap between the currently known upper and lower bounds for approximating $D^*(G)$.

2 Preliminaries

Here, we introduce a formal framework for working with line embeddings and establish two easy “folklore” theorems and some lemmas that will be useful later. The lemmas are formally proved in Appendix A.

2.1 Basic Definitions and Notation

Definition 1. Let (X, d_X) and (Y, d_Y) be metric spaces, with $|X| < \infty$. An *embedding* of (X, d_X) into (Y, d_Y) is an injection $\varphi : X \rightarrow Y$. For any such φ we define

$$D_-(\varphi) := \max_{\substack{u, v \in X \\ u \neq v}} \frac{d_X(u, v)}{d_Y(\varphi(u), \varphi(v))}; \quad D_+(\varphi) := \max_{\substack{u, v \in X \\ u \neq v}} \frac{d_Y(\varphi(u), \varphi(v))}{d_X(u, v)}; \quad D(\varphi) := D_-(\varphi)D_+(\varphi).$$

We call $D_-(\varphi)$ the contraction, $D_+(\varphi)$ the expansion and $D(\varphi)$ the *distortion* of the embedding φ . We say that φ is *non-contracting* if $D_-(\varphi) = 1$. We say that φ is a *line embedding* if $Y = \mathbb{R}$ and $d_Y(u, v) = |u - v|$ for $u, v \in \mathbb{R}$.

The target space \mathbb{R} allows us to rescale any embedding into an equivalent non-contracting one without affecting the distortion. Therefore, we may assume w.l.o.g. that all line embeddings are non-contracting.

Throughout this paper, graphs will be assumed to be simple, undirected and unweighted. For a graph G , we let $V(G)$ denote its vertex set, $E(G)$ its edge set and d_G its shortest path metric, and we define a line embedding of G to be a line embedding of the metric space $(V(G), d_G)$.

Definition 2. The *line distortion* of G , denoted $D^*(G)$, is defined to be $\inf\{D(\varphi) : \varphi \text{ is a non-contracting line embedding of } G\}$.

Any line embedding φ of G naturally induces a total order “ $<_\varphi$ ” on $V(G)$: for $u, v \in V(G)$, we write $u <_\varphi v$ iff $\varphi(u) < \varphi(v)$. A little thought shows that one can also work backwards: any total order \triangleleft on $V(G)$ induces a certain distortion-minimal non-contracting line embedding “ ψ_\triangleleft ” of G , that we call a *canonical embedding*. We define this precisely below and prove its key property (Lemma 4).

Definition 3. Let \triangleleft be a total order on a nonempty finite set S . We use $\min(\triangleleft)$ to denote the minimum element of S according to \triangleleft . If $x \in S$, with $x \neq \min(\triangleleft)$, we use $\text{pred}_\triangleleft(x)$ to denote the predecessor of x according to \triangleleft . For $x, y \in S$, we also define open, closed and half-open intervals according to \triangleleft as follows:

$$\begin{aligned} (x, y)_\triangleleft &:= \{z \in S : x \triangleleft z \triangleleft y\}, & [x, y)_\triangleleft &:= \{z \in S : x \triangleleft z \leq y\}, \\ (x, y]_\triangleleft &:= \{z \in S : x \triangleleft z \leq y\}, & [x, y]_\triangleleft &:= \{z \in S : x \leq z \leq y\}. \end{aligned}$$

Lemma 4 (Canonical Embedding Lemma). *Given a total order \triangleleft on $V(G)$, for a graph G , let $\psi_\triangleleft : V(G) \rightarrow \mathbb{R}$ be defined thus:*

$$\psi_\triangleleft(u) := \begin{cases} 0, & \text{if } u = \min(\triangleleft), \\ \psi_\triangleleft(\text{pred}_\triangleleft(u)) + d_G(\text{pred}_\triangleleft(u), u), & \text{otherwise.} \end{cases}$$

Then ψ_\triangleleft is a non-contracting line embedding of G . Moreover, if φ is any non-contracting line embedding of G with $<_\varphi \equiv \triangleleft$, then $D(\varphi) \geq D(\psi_\triangleleft)$. \square

Definition 5 (Spread). Let φ be a line embedding of (X, d_X) and let $S \subseteq X$. The *spread* of S under φ is defined to be $\max_{v \in S} \varphi(v) - \min_{v \in S} \varphi(v)$.

Definition 6 (Local density). The *local density* $\text{ld}(G)$ of a connected graph G is defined as follows. $\text{ld}(G) := \max\{|B_G(v, \rho)|/\rho : v \in V(G), \rho \geq 1\}$, where $B_G(v, \rho) := \{u \in V(G) : d_G(u, v) \leq \rho\}$.

2.2 Basic Theorems

We now formally prove two simple theorems that provide basic upper and lower bounds on the line distortion of an arbitrary graph.

Theorem 7 (Local density lower bound). *Every connected graph G satisfies $D^*(G) = \Omega(\text{ld}(G))$.*

Proof. Let φ be a non-contracting line embedding of G . Let $v \in V(G)$ and $\rho \geq 1$ be chosen such that $\text{ld}(G) = |S|/\rho$, where $S := B_G(v, \rho)$. Since φ is non-contracting, the spread of S under φ is at least $|S| - 1$, i.e., $\exists x, y \in S$ such that $\varphi(x) - \varphi(y) \geq |S| - 1$. But $d_G(x, y) \leq 2\rho$. Therefore, $D(\varphi) \geq D_+(\varphi) \geq (|S| - 1)/(2\rho) = \Omega(\text{ld}(G))$, which completes the proof. \square

Theorem 8 (General upper bound). *Every connected n -vertex graph G satisfies $D^*(G) \leq 2n - 2 = O(n)$. This upper bound is asymptotically optimal: in fact, there exists an n -vertex tree with line distortion $\Omega(n)$.*

Proof. Let $u_1 \triangleleft u_2 \triangleleft \dots \triangleleft u_n$ be any depth first search ordering of $V(G)$ and let T be the corresponding depth first search tree. Let ω_i denote the walk along T from u_i to u_{i+1} . By the Canonical Embedding Lemma, the spread of $V(G)$ under ψ_\triangleleft equals $\sum_{i=1}^{n-1} d_G(u_i, u_{i+1}) \leq \sum_{i=1}^{n-1} \text{length}(\omega_i) \leq 2(n - 1)$, where the final inequality holds because the concatenated walk $\langle \omega_1, \omega_2, \dots, \omega_{n-1} \rangle$ traverses every edge of G at most twice. Therefore $D^*(G) \leq D(\psi_\triangleleft) \leq 2(n - 1)$.

For the optimality, consider $K_{1, n-1}$, the star on n vertices. Since $\text{ld}(K_{1, n-1}) = n$, Theorem 7 implies $D^*(K_{1, n-1}) = \Omega(n)$. \square

2.3 Some Technical Lemmas

We now prove some technical lemmas about line embeddings that apply to all trees and will be useful when analyzing certain specific families of trees later on.

Lemma 9. *Let φ be a non-contracting line embedding of a tree T . Then $D(\varphi) = \max_{\{u,v\} \in E(T)} |\varphi(u) - \varphi(v)|$.*

Proof. This is immediate from the triangle inequality for d_T . \square

Definition 10 (Connected ordering, Parent). Let T be a tree. A total order \triangleleft on $V(T)$ is called a *connected ordering* of T if, for all $u \in V(T)$, the subgraph of T induced by $[\min(\triangleleft), u]_{\triangleleft}$ is connected. Clearly, for any such \triangleleft and any vertex $u \neq \min(\triangleleft)$, there exists a unique vertex v such that $v \triangleleft u$ and $\{u, v\} \in E(T)$; this v is called the *parent* of u according to \triangleleft and is denoted $\text{par}_{\triangleleft}(u)$.

Lemma 11. *Let \triangleleft be a connected ordering of the tree T . Then*

$$D(\psi_{\triangleleft}) = \max_{\substack{u \in V(T) \\ u \neq \min(\triangleleft)}} \sum_{v \in (\text{par}_{\triangleleft}(u), u]_{\triangleleft}} d_T(\text{pred}_{\triangleleft}(v), v). \quad \square$$

Lemma 12. *If T' is obtained from the tree T by contracting an edge of T , then $D^*(T') \leq 2D^*(T)$.* \square

3 Embedding Complete Binary Trees

Let BIN_h denote the complete binary tree of height h . Throughout this section n will denote the number of vertices of BIN_h . Thus, $n = 2^{h+1} - 1$. We now determine the asymptotic behavior of $D^*(\text{BIN}_h)$.

For the upper bound, we use the Canonical Embedding Lemma: we produce a suitable total order \triangleleft on $V(\text{BIN}_h)$ and analyze $D(\psi_{\triangleleft})$. We assume that the two children of each internal vertex of a binary tree are marked as “left” and “right,” so that we can talk of traversing the tree in symmetric order (also known as “in-order”). For a set $S \subseteq V(T)$, where T is a binary tree, let $\sigma(S)$ denote the sequence obtained by arranging the elements of S in symmetric order.

Lemma 13. *Let \prec denote the total order given by $\sigma(V(\text{BIN}_k))$. Then the spread of $V(\text{BIN}_k)$ under ψ_{\prec} is at most 2^{k+1} .*

Proof. Let s_k be the spread in question. Clearly $s_2 = 2$ and $s_{k+1} = 2s_k + 2k$ for $k \geq 2$. A simple inductive argument shows that $s_k \leq 2^{k+1} - 3k$. \square

To avoid notational clutter we assume that h is a power of 2; it will be clear that our argument generalizes to arbitrary h . Let V_i denote the set of vertices at level i of BIN_h , the root being at level 0 and let $A = V_0 \cup V_1 \cup \dots \cup V_{\log h}$.² Let us also define the sets $B_1, \dots, B_{h/2}$, recursively, as follows:

$$B_i = \{v : v \text{ is a descendant of one of the leftmost } 2^{i+1} \text{ vertices in } (V_{i+\log h} \setminus (B_1 \cup \dots \cup B_{i-1}))\}.$$

Here, a vertex is assumed to be a descendant of itself. It is straightforward to check that the following sequence is a permutation of $V(\text{BIN}_h)$:

$$\langle \sigma(A), \sigma(B_1), \sigma(V_{1+\log h} \setminus B_1), \sigma(B_2), \sigma(V_{2+\log h} \setminus B_2), \dots, \sigma(B_{h/2}), \sigma(V_{h/2+\log h} \setminus B_{h/2}) \rangle. \quad (1)$$

²Here, and for the rest of the paper, we assume that logarithms are to the base 2.

Lemma 14. *Let \triangleleft be the total order on $V(\text{BIN}_h)$ given by the sequence (1). Then $D(\psi_{\triangleleft}) \leq 5 \cdot 2^h / h = O(n / \log n)$.*

Proof. By Lemma 9, it suffices to upper bound the stretch $|\psi_{\triangleleft}(u) - \psi_{\triangleleft}(v)|$ of an arbitrary edge $\{u, v\}$ of BIN_h . Suppose u is the parent of v . If $v \in A$ then, by Lemma 13, the stretch is at most $2^{(\log h)+1} = 2h$. If $u \in B_i$, for some i , then edge is part of a complete binary tree of height at most $h - \log h - 1$. By Lemma 13, the stretch is at most $2^{h-\log h} = 2^h / h$.

For any other edge, we must have $u \triangleleft v$, $u \in V_{i-1+\log h}$ and $v \in V_{i+\log h}$, where $1 \leq i \leq h/2$. The construction of the sequence (1) ensures that $(u, v]_{\triangleleft} \subseteq V_{i-1+\log h} \cup V_{i+\log h} \cup B_{i+\log h}$. By Lemma 13, the spread of V_j under ψ_{\triangleleft} is at most 2^{j+1} . Moreover, $B_{i+\log h}$ consists of 2^{i+1} subtrees, each of height $h - i - \log h$. Therefore, each of these subtrees has spread at most $2^{1+h-i-\log h}$. Finally, the distance in BIN_h between the rightmost vertex of one of these subtrees and the leftmost vertex of the next is at most $2h$. Therefore, these subtrees are embedded at most $2h$ apart from each other. Putting it all together,

$$\begin{aligned} \psi_{\triangleleft}(v) - \psi_{\triangleleft}(u) &\leq 2^{i+\log h} + 2^{1+i+\log h} + 2^{i+1} (2^{1+h-i-\log h} + 2h) \\ &= 7 \cdot 2^i h + 2^{h+2} / h \\ &\leq 5 \cdot 2^h / h, \end{aligned}$$

where the final inequality holds because $i \leq h/2$ and h is large enough. \square

Theorem 15. *For h large enough, $D^*(\text{BIN}_h) = \Theta(2^h / h) = \Theta(n / \log n)$.*

Proof. The upper bound follows from Lemma 14. The lower bound follows from Theorem 7 because $\text{ld}(\text{BIN}_h) = \Theta(n / \log n)$. \square

4 Embedding Fans and Combs

Let $\text{FAN}_{a,b}$ and $\text{COMB}_{a,b}$ denote, respectively, the fan and the comb with a arms, each of length b . For illustrative examples, please see Fig. 1. We assume that the arms of $\text{FAN}_{a,b}$ are numbered 1 to a in some arbitrary order, and that the vertices on each arm are numbered from 0 to b in increasing order of distance from the root of the fan. For $1 \leq i \leq a$ and $0 \leq j \leq b$, let v_{ij} denote the j th vertex on the i th arm of the fan. Note that the root of the fan is $v_{10} = v_{20} = \dots = v_{a0}$ and that

$$E(\text{FAN}_{a,b}) = \{\{v_{ij}, v_{i,j+1}\} : 1 \leq i \leq a, 0 \leq j \leq b-1\}.$$

For $\text{COMB}_{a,b}$, we assume that its vertices are numbered according to an $a \times (b+1)$ grid with w_{ij} denoting the j th vertex on the i th arm of the comb:

$$E(\text{COMB}_{a,b}) = \{\{w_{i0}, w_{i+1,0}\} : 1 \leq i \leq a-1\} \cup \{\{w_{ij}, w_{i,j+1}\} : 1 \leq i \leq a, 0 \leq j \leq b-1\}.$$

We shall determine the asymptotic behavior of $D^*(\text{FAN}_{a,b})$ and $D^*(\text{COMB}_{a,b})$ for large a and b . Our results immediately will imply that $D^*(G) / \text{ld}(G)$ can be as large as $\Omega(n)$, for an n -vertex graph G .

4.1 Upper Bounds

As before, upper bounds follow by analyzing $D(\psi_{\triangleleft})$ for a suitable total order \triangleleft on the vertices of the appropriate graph. The comb is especially simple: the sequence $\langle \langle w_{ij} \rangle_{j=0}^b \rangle_{i=1}^a$ is easily seen to give a connected ordering \triangleleft of $\text{COMB}_{a,b}$. For any vertex $w_{ij} \neq w_{10}$, we have $d_{\text{COMB}}(\text{pred}_{\triangleleft}(w_{ij}), w_{ij}) = b+1$, if $j = 0$, and 1, otherwise. Applying Lemma 11 we get $D(\psi_{\triangleleft}) = 2b+1$. Thus, we have proved:

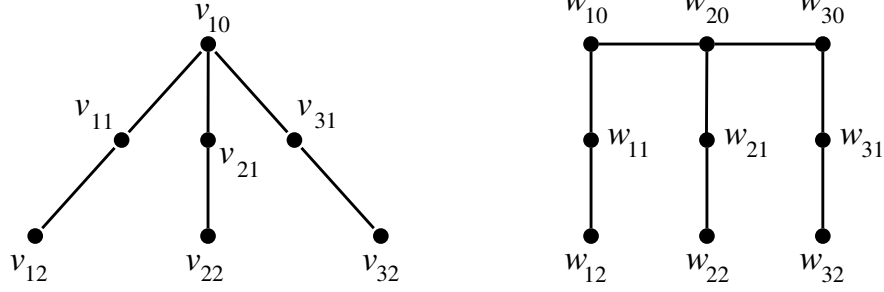


Figure 1: The graphs $\text{FAN}_{3,2}$ and $\text{COMB}_{3,2}$, together with their vertex labels as in the discussion in Section 4.

Theorem 16. $D^*(\text{COMB}_{a,b}) \leq 2b + 1 = O(b)$. □

We now turn to upper bounding $D^*(\text{FAN}_{a,b})$. Note that the embedding used above for $\text{COMB}_{a,b}$ was the canonical embedding for a depth first ordering of its vertices. Using a similar depth first ordering for $\text{FAN}_{a,b}$ gives us an embedding whose distortion is as large as $\Theta(ab) = \Theta(n)$. Using a breadth first ordering also gives a similarly weak result. Below, we give a more complicated construction that proves a nontrivial upper bound.

Theorem 17. $D^*(\text{FAN}_{a,b}) = O(\max\{b, a\sqrt{b}\})$.

Proof. Set $r := \lceil a/\sqrt{b} \rceil$. For $1 \leq k \leq \lceil \sqrt{b} \rceil$, define the sequences σ_k and π_k as follows:

$$\begin{aligned} \sigma_k &:= \langle v_{ij} \rangle_{j=k}^b \Big|_{i=(k-1)r+1}^{kr}, \\ \pi_k &:= \langle v_{ik} \rangle_{i=kr+1}^a. \end{aligned}$$

We assume that sequences are automatically trimmed by removing all undefined symbols v_{ij} . Now define \triangleleft to be the total order on $V(\text{FAN}_{a,b})$ given by the following sequence:

$$\langle v_{10}, \sigma_1, \pi_1, \sigma_2, \pi_2, \dots, \sigma_{\lceil \sqrt{b} \rceil}, \pi_{\lceil \sqrt{b} \rceil} \rangle.$$

We claim that $D(\psi_{\triangleleft}) = O(\max\{b, a\sqrt{b}\})$.

To prove our claim, we observe that ψ_{\triangleleft} is a connected ordering of $T := \text{FAN}_{a,b}$. Therefore, we may use Lemma 11, for which we analyze the quantity $f(v) := d_T(\text{pred}_{\triangleleft}(v), v)$ for all vertices $v \neq v_{10}$.

Each sequence σ_k traverses r paths in $\text{FAN}_{a,b}$, each moving away from the root. Color a vertex v red if $\text{pred}_{\triangleleft}(v)$ is the furthest-from-root vertex on any of these paths. Clearly, $f(v) \leq 2b$ for any red vertex v . If $\text{pred}_{\triangleleft}(v)$ occurs in π_k , for some k , then we see that $f(v) \leq 2k + 1 \leq 3\sqrt{b}$; color v blue in this case. For all other (i.e., uncolored) vertices, $\text{pred}_{\triangleleft}(v) = \text{par}_{\triangleleft}(v)$, whence $f(v) = 1$.

By construction of \triangleleft , each interval $(\text{par}_{\triangleleft}(u), u]_{\triangleleft}$ contains at most r red vertices, at most a blue vertices and at most rb uncolored vertices. Therefore,

$$\sum_{v \in (\text{par}_{\triangleleft}(u), u]_{\triangleleft}} f(v) \leq r \cdot 2b + a \cdot 3\sqrt{b} + rb \cdot 1 \leq 3(rb + a\sqrt{b}).$$

If $a \geq \sqrt{b}$, then $r = \lceil a/\sqrt{b} \rceil \leq 2a/\sqrt{b}$, so the above sum is at most $9a\sqrt{b}$. On the other hand, if $a < \sqrt{b}$, then $r = 1$ and $a\sqrt{b} < b$, so the above sum is at most $6b$. Therefore, the above sum is upper bounded by $O(\max\{b, a\sqrt{b}\})$. Lemma 11 completes the proof. □

4.2 Tight Lower Bounds

We now show that the above upper bound on $D^*(\text{FAN}_{a,b})$ is asymptotically tight. This is our most technically involved proof. We give two different arguments: one that proves an $\Omega(b)$ lower bound and one that proves an $\Omega(a\sqrt{b})$ bound. A tight lower bound on $D^*(\text{COMB}_{a,b})$ follows as a corollary of our argument. Note that the local density lower bound (Theorem 7) cannot give us such strong bounds, because these graphs have local density at most a .

Lemma 18. *Any line embedding φ of $\text{FAN}_{3,b}$ satisfies $D(\varphi) \geq 2b$. Therefore $D^*(\text{FAN}_{3,b}) \geq 2b = \Omega(b)$.*

Proof. Let $T := \text{FAN}_{3,b}$. By the Canonical Embedding Lemma (Lemma 4), we may assume that $\varphi = \psi_{\triangleleft}$ for some total order \triangleleft on $V(\text{FAN}_{3,b})$. Assume w.l.o.g. that $v_{1b} \triangleleft v_{2b} \triangleleft v_{3b}$.

Suppose we have $v_{2b} \triangleleft v_{1j}$ for some j with $0 \leq j < b$. Then $v_{1b} \triangleleft v_{2b} \triangleleft v_{1j}$, whence there must exist an integer k , with $0 \leq k < b$, such that $v_{1,k+1} \triangleleft v_{2b} \triangleleft v_{1k}$. By definition of ψ_{\triangleleft} ,

$$\begin{aligned} \psi_{\triangleleft}(v_{1k}) &= \psi_{\triangleleft}(v_{1,k+1}) + \sum_{u \in (v_{1,k+1}, v_{1k}]_{\triangleleft}} d_T(\text{pred}_{\triangleleft}(u), u) \\ &\geq \psi_{\triangleleft}(v_{1,k+1}) + d_T(v_{1,k+1}, v_{2b}) + d_T(v_{2b}, v_{1k}) \\ &\geq \psi_{\triangleleft}(v_{1,k+1}) + 2b, \end{aligned}$$

where the first inequality follows from the triangle inequality for d_T . Therefore $|\psi_{\triangleleft}(v_{1k}) - \psi_{\triangleleft}(v_{1,k+1})| \geq 2b$ whereas $d_T(v_{1k}, v_{1,k+1}) = 1$. Thus $D(\psi_{\triangleleft}) \geq 2b$.

If $v_{3j} \triangleleft v_{2b}$ for some j with $0 \leq j < b$, then a similar argument gives $D(\psi_{\triangleleft}) \geq 2b$.

Therefore, we are left to consider the case when, for all j such that $0 \leq j \leq b$, we have $v_{1j} \triangleleft v_{2b} \triangleleft v_{3j}$. However, $v_{10} = v_{30}$ is the root of the fan T , so setting $j = 0$ gives us a contradiction in this case. \square

Theorem 19. $D^*(\text{COMB}_{3,b}) \geq b/2 = \Omega(b)$.

Proof. Notice that $\text{FAN}_{3,b}$ can be obtained from $\text{COMB}_{3,b}$ by contracting two edges. Applying Lemma 12 and the above lemma gives us $D^*(\text{COMB}_{3,b}) \geq D^*(\text{FAN}_{3,b})/4 \geq b/2$. \square

We now turn to proving an $\Omega(a\sqrt{b})$ bound on $D^*(\text{FAN}_{a,b})$ for large a and b . We will need to consider induced subgraphs of $\text{FAN}_{a,b}$ that contain the root v_{10} and, for each arm of $\text{FAN}_{a,b}$, either all or none of the non-root vertices in the arm. We call these special subgraphs *subfans* of $\text{FAN}_{a,b}$.

Definition 20 (Segments, Weights, Links, Consecutiveness). Let T be a subfan of $\text{FAN}_{a,b}$ and \triangleleft be an total order on $V(T)$. An interval of \triangleleft that induces a path in $T \setminus \{v_{10}\}$ and is maximal with respect to this property is called a *segment* of T . Clearly, $V(T) \setminus \{v_{10}\}$ is partitioned into segments. We define the *weight* $W(s)$ of a segment s to be $\max\{j : v_{ij} \in s\}$. We call an edge of T a *link* if its endpoints either lie in distinct segments or include the root v_{10} . We call two vertices u and v *consecutive* if there is no vertex x such that either $u \triangleleft x \triangleleft v$ or $v \triangleleft x \triangleleft u$; we define consecutiveness for segments similarly. All of the above definitions are with respect to the particular order \triangleleft .

Lemma 21. *There exists a non-contracting line embedding φ of $\text{FAN}_{a,b}$ such that $D(\varphi) = D^*(\text{FAN}_{a,b})$ with the property that no pair of non-adjacent non-root vertices lying on the same arm of $\text{FAN}_{a,b}$ are consecutive with respect to $<_{\varphi}$.*

Proof. Please see Appendix A. \square

Lemma 22. Suppose φ is an embedding of $\text{FAN}_{a,b}$ with the properties guaranteed by Lemma 21, T is a subfan of $\text{FAN}_{a,b}$ and $\lambda = \{v_{ij}, v_{i,j+1}\}$ is a link of T with respect to $<_\varphi$. Define $S(\lambda)$ to be the set of all segments s of T such that

$$\begin{aligned} \exists x \in s (v_{ij} \leq_\varphi x) \text{ and } \forall x \in s (x <_\varphi v_{i,j+1}), & \text{ if } v_{ij} <_\varphi v_{i,j+1} \\ \forall x \in s (v_{i,j+1} <_\varphi x) \text{ and } \exists x \in s (x \leq_\varphi v_{ij}), & \text{ otherwise.} \end{aligned}$$

Then $D(\varphi) \geq \sum_{s \in S(\lambda)} 2W(s)$.

Proof. Assume w.l.o.g. that $v_{i,j+1} <_\varphi v_{ij}$. Slightly abusing notation, let us extend the total order $<_\varphi$ to the segments of T in the natural way. Let $s_1 <_\varphi s_2 <_\varphi \dots <_\varphi s_t$ be the segments in $S(\lambda)$ and let w_k be the vertex in s_k that is farthest from the root. Then $W(s_k) = d_T(v_{10}, w_k)$. Note that $v_{i,j+1} \notin s_1$ but $v_{ij} \in s_t$.

By the property of φ guaranteed by Lemma 21, no two segments of T are consecutive. Therefore w_{k+1} must be in a different arm of T from w_k , for all k with $1 \leq k < t$. Thus,

$$\varphi(w_{k+1}) - \varphi(w_k) \geq d_T(w_k, w_{k+1}) = d_T(w_k, v_{10}) + d_T(v_{10}, w_{k+1}) = W(s_k) + W(s_{k+1}). \quad (2)$$

Observe that w_1 is not on the i th arm of T and that $v_{ij} = w_t$. Therefore,

$$\varphi(w_1) - \varphi(v_{i,j+1}) \geq d_T(v_{i,j+1}, w_1) > d_T(v_{ij}, w_1) = W(s_1) + W(s_t). \quad (3)$$

Adding (3) and all $k - 1$ inequalities given by (2) together gives us $\varphi(v_{ij}) - \varphi(v_{i,j+1}) \geq \sum_{k=1}^t 2W(s_k)$. To finish the proof, we note that $D(\varphi) \geq |\varphi(v_{ij}) - \varphi(v_{i,j+1})|/d_T(v_{ij}, v_{i,j+1}) = |\varphi(v_{ij}) - \varphi(v_{i,j+1})|$. \square

Lemma 23. $D^*(\text{FAN}_{a,b}) = \Omega(a\sqrt{b})$.

Proof. Let φ be the embedding whose existence is guaranteed by Lemma 21. Throughout this proof, segments and links will be with respect to the total order $<_\varphi$. For each integer m with $1 \leq m \leq \lceil \log b \rceil$,³ let T_m denote the subfan of $\text{FAN}_{a,b}$ consisting of exactly those arms in which the number of segments, excluding the root, lies in $[2^{m-1}, 2^m)$. Some of the subfans T_m may be empty. Let a_m denote the number of arms of T_m and let S_m denote the set of segments of T_m .

Let s_m^- and s_m^+ denote the minimum and maximum segment in S_m , respectively, under the total order $<_\varphi$ (i.e., the leftmost and rightmost segment, respectively). Let L_m be the set of links encountered on the path from a vertex in s_m^- to a vertex in s_m^+ . For each link in L_m we obtain an inequality from Lemma 22. Adding these inequalities together gives

$$|L_m| \cdot D(\varphi) \geq \sum_{s \in S_m \setminus \{s_m^-, s_m^+\}} 2W(s). \quad (4)$$

Consider an arm A of T_m that contains exactly t segments (note that $t \geq 2^{m-1}$). One of these segments must include the vertex in A farthest from the root; this segment has weight b . The smallest possible weights for the other segments are $\{1, 2, \dots, t-1\}$. Therefore,

$$\sum_{s \in A} W(s) \geq b + (1 + 2 + \dots + (t-1)) \geq b + \frac{t^2}{4} \geq b + 2^{2m-4}.$$

Using this fact in (4) gives

$$|L_m| \cdot D(\varphi) \geq 2a_m(b + 2^{2m-4}) - 2W(s_m^-) - 2W(s_m^+).$$

³We remind the reader that logarithms are to the base 2.

Consider any m such that $a_m \geq 4$. Using $W(s_m^-) \leq b$, $W(s_m^+) \leq b$ and $|L_m| \leq 2 \cdot 2^m$ in the above inequality gives us $2^{m+1}D(\varphi) \geq a_m(b + 2^{2m-4})$, whence

$$D(\varphi) \geq a_m \left(\frac{b}{2^{m+1}} + 2^{m-5} \right). \quad (5)$$

Suppose there exists an m such that $a_m \geq \min\{2^{m-2}a/\sqrt{b}, a\sqrt{b}/2^{m+2}\}$. Then, by (5), we would have

$$D(\varphi) \geq \min \left\{ \frac{2^{m-2}a}{\sqrt{b}} \cdot \frac{b}{2^{m+1}}, \frac{a\sqrt{b}}{2^{m+2}} \cdot 2^{m-5} \right\} = \frac{a\sqrt{b}}{2^7}$$

and we would be done. Suppose, instead, that no such m exists. Then

$$\sum_{m=1}^{\lceil \log b \rceil} a_m < \sum_{m=1}^{\lceil \log b \rceil} \min \left\{ \frac{2^{m-2}a}{\sqrt{b}}, \frac{a\sqrt{b}}{2^{m+2}} \right\} = \frac{a}{4} \sum_{m=1}^{\lceil \log b \rceil} \min \left\{ \frac{2^m}{\sqrt{b}}, \frac{\sqrt{b}}{2^m} \right\} \leq \frac{a}{4} \sum_{m=-\infty}^{\infty} \min \left\{ \frac{2^m}{\sqrt{b}}, \frac{\sqrt{b}}{2^m} \right\} \leq a$$

and we have a contradiction since, by the definition of a_m , we must have $\sum_{m=1}^{\lceil \log n \rceil} a_m = a$. This completes the proof. \square

4.3 Summary of Results and Consequences

To summarize our results above, we combine Theorems 16, 17 and 19 and Lemmas 18 and 23 to obtain:

Theorem 24. $D^*(\text{COMB}_{a,b}) = \Theta(b)$ and $D^*(\text{FAN}_{a,b}) = \Theta(\max\{b, a\sqrt{b}\})$. \square

Notice that $\text{ld}(\text{FAN}_{a,b}) \leq a$ and $\text{ld}(\text{COMB}_{a,b}) \leq a$. Thus, these two families of trees exhibit arbitrarily large gaps between the local density $\text{ld}(T)$ and the line distortion $D^*(T)$, for appropriately chosen parameters a and b . Indeed, this gap could be made as large as $\Theta(n)$ by considering, e.g., $\text{COMB}_{3,n/3}$. This large gap is to be contrasted with the situation for the *bandwidth* $\text{bw}(T)$: as mentioned in Section 1.1, it has been proven that $\text{bw}(T) = O(\text{ld}(T) \log^{2.5} n)$ for a tree T and that $\text{bw}(G) = O(\text{ld}(G) \log^{3.5} n)$ for a general graph G , and it is conjectured that $\text{bw}(G) = O(\text{ld}(G) \log n)$.

From our results one can see that $D^*(T)$, for an n -vertex tree T , can be arbitrarily low or high, even for the highly restricted class of binary trees (equivalently, trees with maximum degree 3), as shown by the following theorem.

Theorem 25. *Let $f(n)$ be a nondecreasing function of n with $1 \leq f(n) \leq n$. There is a family of binary trees $\{T_i\}$, where T_n has n vertices, such that $D^*(T_n) = \Theta(f(n))$.*

Proof. Define $T_n = \text{COMB}_{\lceil n/f(n) \rceil, \lceil f(n) \rceil}$. When appropriately rooted, T_n is a binary tree. By Theorem 24, $D^*(T_n) = \Theta(f(n))$. \square

5 Concluding Remarks

We have studied bi-Lipschitz embeddings of graph metrics into the real line and determined the optimal distortion of a line embedding of certain special families of trees: namely, complete binary trees, combs and fans. These families illustrate some of the nontrivial combinatorial techniques that arise in either upper or lower bounding their line distortions.

While it would have been nice to have a general theorem applicable to all trees, such as one that relates the optimal distortion to a more well-studied combinatorial notion, we believe that there is no easy theorem of the sort.⁴ Theorem 25 highlights this point. It also shows that our sublinear distortion line embedding of the complete binary tree in Theorem 15 does not generalize to arbitrary binary trees.

Acknowledgments

We would like to thank Anupam Gupta for getting us interested in the problems handled here and Bernard Chazelle and an anonymous referee for helpful comments on earlier drafts of the paper.

References

- [1] S. Arora, S. Rao, and U. V. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *Proc. 36th Annual ACM Symposium on the Theory of Computing*, pages 222–231, 2004.
- [2] M. W. Bern, H. J. Karloff, P. Raghavan, and B. Schieber. Fast geometric approximation techniques and geometric embedding problems. *Theor. Comput. Sci.*, 106(2):265–281, 1992. Preliminary version in *Proc. 5th ACM Symposium on Computational Geometry*, pages 292–301, 1989.
- [3] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math.*, 52:46–52, 1985.
- [4] J. Bourgain. The metrical interpretation of superreflexivity in Banach spaces. *Israel J. Math.*, 56(2):222–230, 1986.
- [5] M. Bădoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Low-distortion embeddings of general metrics into the line. In *Proc. 37th Annual ACM Symposium on the Theory of Computing*, pages 225–233, 2005.
- [6] M. Bădoiu, K. Dhamdhere, A. Gupta, Y. Rabinovich, H. Räcke, R. Ravi, and A. Sidiropoulos. Approximation algorithms for low-distortion embeddings into low-dimensional spaces. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 119–128, 2005.
- [7] K. Diks. Embeddings of binary trees in lines. *Theor. Comput. Sci.*, 36:319–331, 1985.
- [8] U. Feige. Approximating the bandwidth via volume respecting embeddings. *J. Comput. Syst. Sci.*, 60(3):510–539, 2000. Preliminary version in *Proc. 30th Annual ACM Symposium on the Theory of Computing*, pages 90–99, 1998.
- [9] A. Gupta. Embedding trees into low dimensional Euclidean spaces. *Disc. Comput. Geom.*, 24(1):106–116, 2000. Preliminary version in *Proc. 31st Annual ACM Symposium on the Theory of Computing*, pages 694–700, 1999.
- [10] A. Gupta. Improved bandwidth approximation for trees and chordal graphs. *J. Alg.*, 40(1):24–36, 2001. Preliminary version in *Proc. 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 788–793, 2000.
- [11] R. Heckmann, R. Klasing, B. Monien, and W. Unger. Optimal embedding of complete binary trees into lines and grids. *J. Parallel Distrib. Comput.*, 49(1):40–56, 1998.
- [12] P. Indyk. Algorithmic applications of low-distortion geometric embeddings. In *Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 10–33, 2001.
- [13] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proc. 30th Annual ACM Symposium on the Theory of Computing*, pages 604–613, 1998.
- [14] C. Kenyon, Y. Rabani, and A. Sinclair. Low distortion maps between point sets. In *Proc. 36th Annual ACM Symposium on the Theory of Computing*, pages 272–280, 2004.

⁴It is possible that more sophisticated combinatorial notions could capture $D^*(T)$; this is hinted at by a recent result of Lee, Naor and Peres [16], who characterize the optimal distortion of a tree into Hilbert space in terms of edge colorings of the tree.

- [15] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics. In *Proc. 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 434–443, 2004.
- [16] J. R. Lee, A. Naor, and Y. Peres. Trees and markov convexity. *Geom. Funct. Anal.*, 18(5):1609–1659, 2008. Preliminary version in *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1028–1037, 2006.
- [17] T. Leighton and S. Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *Proc. 29th Annual IEEE Symposium on Foundations of Computer Science*, pages 422–431, 1988.
- [18] N. Linial. Finite metric spaces — combinatorics, geometry and algorithms. In *Proc. Intl. Cong. Math. III*, pages 573–586, 2002.
- [19] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. Preliminary version in *Proc. 35th Annual IEEE Symposium on Foundations of Computer Science*, pages 577–591, 1994.
- [20] J. Matoušek. Bi-Lipschitz embeddings into low-dimensional Euclidean spaces. *Comment. Math. Univ. Carolinae*, 31:589–600, 1990.
- [21] J. Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag, 2002.

A Proofs of Various Lemmas

Restatement of Lemma 4 (Canonical Embedding Lemma). Given a total order \triangleleft on $V(G)$, for a graph G , let $\psi_{\triangleleft} : V(G) \rightarrow \mathbb{R}$ be defined thus:

$$\psi_{\triangleleft}(u) := \begin{cases} 0, & \text{if } u = \min(\triangleleft), \\ \psi_{\triangleleft}(\text{pred}_{\triangleleft}(u)) + d_G(\text{pred}_{\triangleleft}(u), u), & \text{otherwise.} \end{cases}$$

Then ψ_{\triangleleft} is a non-contracting line embedding of G . Moreover, if φ is any non-contracting line embedding of G with $<_{\varphi} \equiv \triangleleft$, then $D(\varphi) \geq D(\psi_{\triangleleft})$.

Proof. We prove that ψ_{\triangleleft} is non-contracting by contradiction. Suppose there exist distinct $u, v \in V(G)$ such that $u \triangleleft v$ and $\psi_{\triangleleft}(v) - \psi_{\triangleleft}(u) < d_G(u, v)$. Moreover, suppose u and v are as close as possible in the total order \triangleleft . If there exists $x \in V(G)$ with $u \triangleleft x \triangleleft v$, then, by the triangle inequality in G ,

$$(\psi_{\triangleleft}(v) - \psi_{\triangleleft}(x)) + (\psi_{\triangleleft}(x) - \psi_{\triangleleft}(u)) \geq d_G(v, x) + d_G(x, u) \geq d_G(v, u),$$

which gives us a contradiction. On the other hand, if $u = \text{pred}_{\triangleleft}(v)$, then the definition of ψ_{\triangleleft} gives us a contradiction.

Now suppose φ is any non-contracting line embedding of G with $<_{\varphi} \equiv \triangleleft$. Let $u_1 \triangleleft u_2 \triangleleft \dots \triangleleft u_n$ be the maximal chain of \triangleleft . Then, for all $i \in \{1, \dots, n-1\}$,

$$\varphi(u_{i+1}) - \varphi(u_i) \geq d_G(u_i, u_{i+1}) = \psi_{\triangleleft}(u_{i+1}) - \psi_{\triangleleft}(u_i),$$

where the first inequality holds because $u_i <_{\varphi} u_{i+1}$ and φ is non-contracting. Thus, for $1 \leq i < j \leq n$,

$$\varphi(u_j) - \varphi(u_i) = \sum_{h=i}^{j-1} (\varphi(u_{h+1}) - \varphi(u_h)) \geq \sum_{h=i}^{j-1} (\psi_{\triangleleft}(u_{h+1}) - \psi_{\triangleleft}(u_h)) = \psi_{\triangleleft}(u_j) - \psi_{\triangleleft}(u_i),$$

which implies $D(\varphi) \geq D(\psi_{\triangleleft})$. □

Restatement of Lemma 11. Let \triangleleft be a connected ordering of the tree T . Then

$$D(\psi_{\triangleleft}) = \max_{\substack{u \in V(T) \\ u \neq \min(\triangleleft)}} \sum_{v \in (\text{par}_{\triangleleft}(u), u]_{\triangleleft}} d_T(\text{pred}_{\triangleleft}(v), v).$$

Proof. Since \triangleleft is a connected ordering, $E(T) = \{\{\text{par}_{\triangleleft}(u), u\} : u \in V(T) \setminus \{\min(\triangleleft)\}\}$. By Lemma 9,

$$D(\psi_{\triangleleft}) = \max_{u \in V(T) \setminus \{\min(\triangleleft)\}} \{\psi_{\triangleleft}(u) - \psi_{\triangleleft}(\text{par}_{\triangleleft}(u))\}.$$

To finish the proof, we note that

$$\begin{aligned} \psi_{\triangleleft}(u) - \psi_{\triangleleft}(\text{par}_{\triangleleft}(u)) &= \sum_{v \in (\min(\triangleleft), u]_{\triangleleft}} d_T(\text{pred}_{\triangleleft}(v), v) - \sum_{v \in (\min(\triangleleft), \text{par}_{\triangleleft}(u)]_{\triangleleft}} d_T(\text{pred}_{\triangleleft}(v), v) \\ &= \sum_{v \in (\text{par}_{\triangleleft}(u), u]_{\triangleleft}} d_T(\text{pred}_{\triangleleft}(v), v). \end{aligned} \quad \square$$

Restatement of Lemma 12. If the tree T' is obtained from the tree T by contracting an edge of T , then $D^*(T') \leq 2D^*(T)$.

Proof. Let φ be any non-contracting line embedding of T and let φ' denote φ restricted to $V(T')$. For any two distinct vertices $u, v \in V(T')$, we have $d_{T'}(u, v) \in \{d_T(u, v), d_T(u, v) - 1\}$. This implies that φ' is a non-contracting line embedding of T' and that $d_{T'}(u, v) \geq d_T(u, v)/2$. Therefore

$$D(\varphi') = \max_{\substack{u, v \in V(T') \\ u \neq v}} \frac{|\varphi'(u) - \varphi'(v)|}{d_{T'}(u, v)} \leq \max_{\substack{u, v \in V(T) \\ u \neq v}} \frac{2 \cdot |\varphi(u) - \varphi(v)|}{d_T(u, v)} = D(\varphi).$$

Picking φ such that $D(\varphi) = D^*(T)$ completes the proof. \square

Restatement of Lemma 21. There exists a non-contracting line embedding φ of $\text{FAN}_{a,b}$ such that $D(\varphi) = D^*(\text{FAN}_{a,b})$ with the property that no pair of non-adjacent non-root vertices lying on the same arm of $\text{FAN}_{a,b}$ are consecutive with respect to $<_{\varphi}$.

Proof. Consider any non-contracting line embedding φ of $T := \text{FAN}_{a,b}$. Suppose v_{ij} and v_{ik} (on the i th arm of T) are a “bad pair,” i.e., they are consecutive with respect to $<_{\varphi}$ but are also non-adjacent non-root vertices: $0 < j < k - 1$. We describe an operation on φ that will eliminate this bad pair. In the sequel we shall apply this operation repeatedly.

Assume w.l.o.g. that $v_{ij} <_{\varphi} v_{ik}$. Define the function $\varphi' : V(T) \rightarrow \mathbb{R}$ to be identical to φ except at vertices v_{il} with $j < l < k$, which are moved so that they lie between $\varphi(v_{ij})$ and $\varphi(v_{ik})$ and are evenly spaced out. To be precise,

$$\varphi'(v_{ml}) = \begin{cases} \varphi(v_{ij}) + (\varphi(v_{ik}) - \varphi(v_{ij})) \cdot \frac{l-j}{k-j}, & \text{if } m = i \text{ and } j < l < k, \\ \varphi(v_{ml}), & \text{otherwise.} \end{cases}$$

We claim that φ' is a non-contracting line embedding of T with $D(\varphi') \leq D(\varphi)$. We say that a vertex v has *moved* if $\varphi'(v) \neq \varphi(v)$. Since φ is non-contracting, we have

$$1 \leq \alpha := \frac{|\varphi(v_{ik}) - \varphi(v_{ij})|}{k-j} \leq D(\varphi).$$

For any pair of distinct vertices $x, y \in V(T)$, three cases arise. If neither x nor y has moved, then $|\varphi'(x) - \varphi'(y)| = |\varphi(x) - \varphi(y)|$. If they have both moved, then, by construction, $|\varphi'(x) - \varphi'(y)| = \alpha d_T(x, y)$. If $x = v_{il}$ has moved and y has not, assume that $y <_\varphi v_{ij}$ (we can make a similar argument if $v_{ik} <_\varphi y$). Then

$$|\varphi'(x) - \varphi'(y)| = |\varphi'(x) - \varphi(v_{ij})| + |\varphi(v_{ij}) - \varphi(y)| \geq d_T(x, v_{ij}) + d_T(v_{ij}, y) \geq d_T(x, y).$$

Furthermore, the path in T from x to y must pass either through v_{ij} or through v_{ik} ; let v denote the vertex among these two that it does pass through. Then

$$|\varphi'(x) - \varphi'(y)| \leq |\varphi'(x) - \varphi(v)| + |\varphi(v) - \varphi(y)| \leq \alpha d_T(x, v) + D(\varphi) d_T(v, y) \leq D(\varphi) d_T(x, y).$$

Thus, in all three cases, $d_T(x, y) \leq |\varphi'(x) - \varphi'(y)| \leq D(\varphi) d_T(x, y)$, which proves our claim.

By repeatedly applying the above operation, we can eliminate all bad pairs without ever increasing the distortion. To see that this procedure eventually terminates, observe that the operation strictly decreases the number of segments of T with respect to $<_\varphi$. \square