

Constant Approximation Algorithms for Embedding Graph Metrics into Trees and Outerplanar Graphs

V. Chepoi · F.F. Dragan · I. Newman ·
Y. Rabinovich · Y. Vaxès

Received: 3 July 2010 / Revised: 27 October 2011 / Accepted: 2 November 2011 /

Published online: 23 November 2011

© Springer Science+Business Media, LLC 2011

Abstract In this paper, we present a simple factor 6 algorithm for approximating the optimal multiplicative distortion of embedding a graph metric into a tree metric (thus improving and simplifying the factor 100 and 27 algorithms of Bădoiu et al. (Proceedings of the eighteenth annual ACM–SIAM symposium on discrete algorithms (SODA 2007), pp. 512–521, 2007) and Bădoiu et al. (Proceedings of the 11th international workshop on approximation algorithms for combinatorial optimization problems (APPROX 2008), Springer, Berlin, pp. 21–34, 2008)). We also present a constant factor algorithm for approximating the optimal distortion of embedding a graph metric into an outerplanar metric. For this, we introduce a general notion of a metric relaxed minor and show that if G contains an α -metric relaxed H -minor, then the distortion of any embedding of G into any metric induced by a H -minor

An extended abstract of this paper appeared in the proceedings of APPROX-RANDOM 2010.

V. Chepoi (✉) · Y. Vaxès

Laboratoire d'Informatique Fondamentale, Faculté des Sciences de Luminy, Université d'Aix-Marseille, 13288 Marseille Cedex 9, France

e-mail: chepoi@lif.univ-mrs.fr

Y. Vaxès

e-mail: vaxes@lif.univ-mrs.fr

F.F. Dragan

Computer Science Department, Kent State University, Kent, OH 44242, USA

e-mail: dragan@cs.kent.edu

I. Newman · Y. Rabinovich

Department of Computer Science, University of Haifa, Mount Carmel, Haifa 31905, Israel

I. Newman

e-mail: ilan@cs.haifa.ac.il

Y. Rabinovich

e-mail: yuri@cs.haifa.ac.il

free graph is $\geq \alpha$. Then, for $H = K_{2,3}$, we present an algorithm which either finds an α -relaxed minor, or produces an $O(\alpha)$ -embedding into an outerplanar metric.

Keywords Metric embedding · Distortion · Trees · Outerplanar graphs · Approximation algorithms

1 Introduction

1.1 Avant-Propos

The structure of the shortest-path metrics of special classes of graphs, in particular, graph families defined by a set of forbidden minors (e.g., line metrics, tree metrics, planar metrics) is one of the main areas in the theory of metric spaces. From the algorithmic point of view, such metrics typically have more structure than general metrics, and this structure can often be exploited algorithmically. Thus, if the input metric can be well approximated by a special metric, this usually leads to an algorithmic advantage; see, e.g., [14] for a survey of embeddings and their algorithmic applications. One way of understanding this structure is to study the low distortion embeddings from one metric class to another. To do this successfully, one needs to develop tools allowing a decomposition of the host space consistent with the embedded space. If this is impossible, one usually learns much about the limitations of the host space and the richness of the embedded space. In this paper, we pursue this direction and study the embeddings into tree metrics and the metrics of $K_{2,3}$ -minor free graphs (the outerplanar metrics essentially because each 2-connected component of a $K_{2,3}$ -minor free graph is either outerplanar or a K_4).

The study of tree metrics can be traced back to the beginning of the twentieth century, when it was first realized that weighted trees can in some cases serve as an (approximate) model for the description of evolving systems. More recently, as indicated in [17], it was observed that certain Internet originated metrics display tree-like properties. It is well known [18] that tree metrics have a simple structure: d is a tree metric if and only if all submetrics of d of size 4 are such. Moreover, the underlying tree is unique, easily reconstructible, and has rigid local structure corresponding to the local structure of d . But what about the structure of *approximately* tree metrics? We have only partial answers for this question, and yet what we already know seems to indicate that a rich theory might well be hiding there. The strongest results were obtained, so far, for the *additive* distortion. A research on the algorithmic aspects of finding a tree metric of least additive distortion has culminated in [1] (see also [8]), where a 6-approximation algorithm was established (in the notation of [1], it is a 3-approximation algorithm, however, in our more restrictive definition, requiring that the metric is dominated by the approximating one, it is a 6-approximation), together with a (rather close) hardness result. Relaxing the local condition on d by allowing its size-4 submetrics to be δ -close to a tree metric, one gets precisely Gromov's δ -hyperbolic geometry. For a study of algorithmic and other aspects of such geometries, see, e.g., [7, 15].

The situation with the *multiplicative* distortion is less satisfactory. The best result for embedding general metrics into tree metrics is obtained in [4]: the approximation

factor is exponential in $\sqrt{\log \Delta} / \log \log n$, where Δ is the aspect ratio. Judging from the parallel results of [2] for embedding into line metrics, it is conceivable that any constant factor approximation for optimal embedding general metrics into tree metrics is NP-hard. For some small constant γ , the hardness result of [1] implies that it is NP-hard to approximate the multiplicative distortion better than γ even for metrics that come from unit-weighted graphs. For a special interesting case of shortest path metrics of *unit-weighted* graphs, [4] gets a large (around 100) constant approximation factor (which was improved in [3] to a factor of 27). The proof introduces a certain metric-topological obstacle for getting embeddings of distortion better than α , and then algorithmically either produces an $O(\alpha)$ -embedding, or an α -obstacle. Let us mention that such an obstacle was used also in [12], and, essentially, in [16].

1.2 Our Results

In this paper, we study the embeddings of *unweighted* (i.e., unit-weighted) graph metrics into tree metrics and outerplanar metrics. Using a decomposition procedure developed earlier in [5, 6], we simplify and improve the construction of [4] for embedding into tree metrics. The improved constant is 6. We also introduce the notions of relaxed and metric relaxed minors and show that if G contains an α -metric relaxed H -minor, then the distortion of any embedding of the metric of G into any metric induced by a H -minor free graph is at least α . This generalizes the obstacle of [4]. Using this newly defined H -obstacle, we are able to show that it is an essential obstacle not only for trees, but also for graphs without $H = K_{2,3}$ minors as well. We further develop an efficient algorithm which either embeds the input graph G into an outerplanar metric with distortion $O(\alpha)$, or finds an α -metric relaxed $K_{2,3}$ -minor in G . This is a first result of this kind for any H different from a C_4 (which is the corresponding α -metric relaxed minor corresponding to the four-point condition used for embedding into tree-metrics; α -metric relaxed minors are defined in Sect. 4.2). It is our feeling that this obstacle may prove essential for other forbidden H 's, notably $K_{2,r}$, hopefully series-parallel graphs, and beyond. Note also that our proof for embedding into outerplanar graphs has an interesting consequence: it implies that if a graph embeds with distortion α into an outerplanar graph then it embeds with distortion $O(\alpha)$ into a “tree of cycles”, i.e., into an outerplanar graph in which every 2-connected component is either an edge, or a cycle. In particular, this shows that every unweighted outerplanar graph embeds into a tree of cycles with constant distortion.

1.3 Preliminaries

A metric space (X, d) is *isometrically embeddable* into a host metric space (Y, d') if there exists a map $\varphi : X \mapsto Y$ such that $d'(\varphi(x), \varphi(y)) = d(x, y)$ for all $x, y \in X$. In this case, we say that X is a subspace of Y . More generally, $\varphi : X \mapsto Y$ is an *embedding with (multiplicative) distortion $\lambda \geq 1$* if $d(x, y) \leq d'(\varphi(x), \varphi(y)) \leq \lambda \cdot d(x, y)$ for all $x, y \in X$ (note that embedding here is non-contracting; this could be relaxed to contracting embeddings as well, but will not be important in what follows). Given a metric space (X, d) and a class \mathcal{M} of host metric spaces, we denote by

$\lambda^* := \lambda^*(X, \mathcal{M})$ the minimum distortion of an embedding of (X, d) into a member of \mathcal{M} . Analogously, $\varphi : X \mapsto Y$ is an *embedding with additive distortion* $\lambda \geq 0$ if $d(x, y) \leq d'(\varphi(x), \varphi(y)) \leq d(x, y) + \lambda$ for all $x, y \in X$. In a similar way, we can define the minimum additive distortion for embedding of a metric space (X, d) into a class \mathcal{M} of host metric spaces. In this paper, we consider finite connected unweighted graphs as input metric spaces and *tree metrics* (trees) or *outerplanar metrics* (and they relatives) as the class of host metric spaces. If not specified, all our results concern embeddings with multiplicative distortion. For a connected unweighted graph $G = (V, E)$, we denote by $d_G(u, v)$ the shortest-path distance between two vertices u and v of G . A finite metric space (X, d) is called a *tree metric* if it isometrically embeds into a tree, i.e., there exists a weighted tree $T = (X', E')$ such that $X \subseteq X'$ and $d(u, v) = d_T(u, v)$ for any two points $u, v \in X$, where $d_T(u, v)$ is the length of the unique path connecting u and v in T . Analogously, an *outerplanar metric* is a metric space isometrically embeddable into an outerplanar weighted graph. We denote by \mathcal{T} the class of all tree metric spaces and by \mathcal{O} the class of outerplanar metric spaces. Note that \mathcal{T} is a proper subclass of \mathcal{O} .

2 Preliminary Results

In this section, we establish some properties of layering partitions and of embeddings with distortion λ of graph metrics into weighted graphs.

2.1 Layering Partitions

We now briefly describe the layering partitions and establish some of their properties. The layering partitions have been introduced in [5, 6] and recently used in a slightly more general forms in both approximation algorithms for embedding graph metric into trees [3, 4] as well as in some other similar contexts [7, 10, 11].

Let $G = (V, E)$ be an unweighted connected graph with a distinguished vertex s and let $r := \max\{d_G(s, x) : x \in V\}$. A *layering* of G with respect to s is the decomposition of V into the *spheres* $L^i = \{u \in V : d(s, u) = i\}$, $i = 0, 1, 2, \dots, r$. A *layering partition* $\mathcal{LP}(s) = \{L_1^i, \dots, L_{p_i}^i : i = 0, 1, 2, \dots, r\}$ of G is a partition of each L^i into *clusters* $L_1^i, \dots, L_{p_i}^i$ such that two vertices $u, v \in L^i$ belong to the same cluster L_j^i if and only if they can be connected by a path outside the ball $B_{i-1}(s)$ of radius $i - 1$ centered at s . It was shown in [6] that for a given unweighted graph G such a layering partition can be found in linear time. Let Γ be a graph whose vertex set is the set of all clusters L_j^i in a layering partition \mathcal{LP} of a graph G . Two vertices $C = L_j^i$ and $C' = L_{j'}^{i'}$ are adjacent in Γ if and only if there exist $u \in L_j^i$ and $v \in L_{j'}^{i'}$ such that u and v are adjacent in G (see Fig. 1). It is shown in [6] that Γ is a tree, called the *layering tree* of G , and that Γ is computable in linear time in the size of G . In what follows, we assume that Γ is rooted at cluster $\{s\}$.

We can construct a new tree $H = (V, F)$ for a graph G (closely reproducing the global structure of the layering tree Γ) by identifying for each cluster $C = L_j^i \in \mathcal{LP}$ an arbitrary vertex $x_C \in L^{i-1}$ which has a neighbor in $C = L_j^i$ and by making x_C adjacent in H with all vertices $v \in C$ (see the rightmost picture in Fig. 1). Vertex x_C

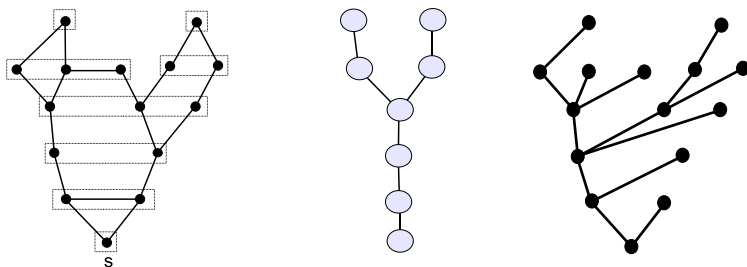


Fig. 1 A layering partition of G and the trees Γ and H associated with this layering partition

will be called *support vertex* for cluster $C = L_j^i$. In what follows, we assume that H is rooted at vertex s .

Let D be the largest diameter of a cluster in a layering partition \mathcal{LP} of G , i.e., $D := \max_{C \in \mathcal{LP}} \max_{v, u \in C} \{d_G(u, v)\}$. Then, the following result (also implicitly used in [5–7] in particular cases) shows that the additive distortion of the embedding of G into H is essentially D :

Proposition 1 *For any vertices x, y of G , $d_H(x, y) - 2 \leq d_G(x, y) \leq d_H(x, y) + D$.*

Proof Let C_x and C_y be the clusters containing the vertices x and y , respectively. Denote by C the cluster which is the nearest common ancestor of C_x and C_y in the layering tree Γ of G . The assertion is trivial if $C_x = C_y = C$. For $C \neq C_x$, let $x', y' \in C$ be the ancestors of x and y , respectively, in a BFS(G, s)-tree. Then $d_\Gamma(C_x, C) = d_G(x, x')$ and $d_\Gamma(C_y, C) = d_G(y, y')$. By construction of H , $d_H(x, y)$ is equal either to $d_\Gamma(C_x, C) + d_\Gamma(C_y, C)$ or to $d_\Gamma(C_x, C) + d_\Gamma(C_y, C) + 2$. Thus, by the triangle inequality,

$$\begin{aligned} d_G(x, y) &\leq d_G(x, x') + d_G(x', y') + d_G(y, y') \leq d_\Gamma(C_x, C) + d_\Gamma(C_y, C) + D \\ &\leq d_H(x, y) + D. \end{aligned}$$

On the other hand, by definition of clusters, $d_G(x, y) \geq d_G(x, x') + d_G(y, y') \geq d_H(x, y) - 2$. \square

Note that tree H , like any BFS-tree, preserves graph distances between the root s and any other vertex of G . We can locally modify H by assigning uniform weights to its edges or by adding Steiner points to obtain a number of other desired properties (like, non-expansiveness, non-contractibility, etc.). For example, assigning length $w := D + 1$ to each edge of H , we will get a uniformly weighted tree $H_w = (V, F, w)$ in which G embeds with multiplicative distortion essentially equal to $D + 1$:

Corollary 1 *For any vertices u, v of G , $d_G(u, v) \leq d_{H_w}(u, v) \leq (D + 1)(d_G(u, v) + 2)$.*

By adding Steiner points and using edge lengths 0 and 1, the tree H can be easily transformed into a tree H' which has the same additive distortion and satisfies the

non-expansive property. For this, for each cluster $C := L_j^i$ we introduce a Steiner point p_C , and add an edge of length 0 between any vertex of C and p_C and an edge of length 1 between p_C and the support vertex x_C for C , defined above.

Corollary 2 *For any vertices u, v of G , $d_{H'}(u, v) \leq d_G(u, v) \leq d_{H'}(u, v) + D$.*

By replacing each edge in H' with edge of length $w := \frac{D+1}{2}$, we obtain a tree H'_w with the following property:

Corollary 3 *For any vertices u, v of G , $d_G(u, v) \leq d_{H'_w}(u, v) \leq (D+1)(d_G(u, v) + 1)$.*

2.2 Embeddings with Distortion λ of Graph Metrics

We continue with two auxiliary standard results about embeddings.

Lemma 1 *If $G = (V, E)$ is an unweighted graph, $G' = (V', E')$ is a weighted graph, and $\varphi : V \mapsto V'$ is a mapping such that $d_{G'}(\varphi(u), \varphi(v)) \leq \lambda$ for any edge uv of G , then $d_{G'}(\varphi(x), \varphi(y)) \leq \lambda d_G(x, y)$ for any pair of vertices x, y of G .*

Proof Consider a shortest path P of G between arbitrary vertices x, y of G . For each edge uv of P , $\varphi(u)$ and $\varphi(v)$ are connected in G' by a path P_{uv} of length $\leq \lambda$. Hence, $\varphi(x)$ and $\varphi(y)$ can be connected in the subgraph of G' induced by $\cup\{P_{uv} : uv \text{ is an edge of } P\}$ by a path with total length of edges at most $\lambda d_G(x, y)$. Hence, $d_{G'}(\varphi(x), \varphi(y)) \leq \lambda d_G(x, y)$. \square

Lemma 2 *If $G = (V, E)$ is an unweighted graph, $G' = (V', E')$ is a weighted graph, and $\varphi : V \mapsto V'$ is a mapping such that $d_{G'}(\varphi(u), \varphi(v)) \geq d_G(u, v)$ for any edge uv of G , then $d_{G'}(\varphi(x), \varphi(y)) \geq d_G(x, y)$ for any pair of vertices x, y of G .*

Proof We proceed by induction on the number of edges in a shortest path between $\varphi(x)$ and $\varphi(y)$ in G' . If $\varphi(x)$ and $\varphi(y)$ are adjacent in G' , then we are done by our condition. Otherwise, let $\varphi(x')$ be the neighbor of $\varphi(x)$ in such a shortest path. By induction hypothesis, $d_{G'}(\varphi(x'), \varphi(y)) \geq d_G(x', y)$. Since $d_{G'}(\varphi(x), \varphi(x')) \geq d_G(x, x')$, the triangle inequality yields $d_{G'}(\varphi(x), \varphi(y)) = d_{G'}(\varphi(x), \varphi(x')) + d_{G'}(\varphi(x'), \varphi(y)) \geq d_G(x, x') + d_G(x', y) \geq d_G(x, y)$, and we are done. \square

3 Approximation Algorithm for Embedding Graph Metrics into Trees

We describe now a simple factor 6 algorithm for approximating the optimal distortion $\lambda^* = \lambda^*(G, T)$ of embedding finite unweighted graphs G into trees. For this, we first investigate the properties of layering partitions of graphs which λ -embed into trees, i.e., for each such graph $G = (V, E)$ there exists a tree $T = (V', E')$ with $V \subseteq V'$ such that

$$d_G(x, y) \leq d_T(x, y) \quad (\text{non-contractibility}) \quad (1)$$

and

$$d_T(x, y) \leq \lambda \cdot d_G(x, y) \quad (\text{bounded expansion}) \quad (2)$$

for every $x, y \in V$.

Denote by $P_T(x, y)$ the path connecting the vertices x, y in T . For $x \in V'$ and $A \subseteq V'$, we denote by $d_T(x, A) = \min\{d_T(x, v) : v \in A\}$ the distance from x to A . First, we show that the diameters of clusters in a layering partition of such a graph G are at most 3λ , allowing already to build a tree with distortion $8\lambda^*$. Refining this property of layering partitions, we construct in $O(|V||E|)$ time a tree into which G embeds with distortion $\leq 6\lambda^*$.

Lemma 3 *If a graph G λ -embeds into a tree, then for any $x, y \in V$, any path $P_G(x, y)$ of G between x, y and any vertex $c \in P_T(x, y)$, we have $d_T(c, P_G(x, y)) \leq \lambda/2$.*

Proof Removing c from T , we separate x from y . Let T_y be the subtree of $T \setminus \{c\}$ containing y . Since $x \notin T_y$, we can find an edge ab of $P_G(x, y)$ with $a \in T_y$ and $b \notin T_y$. Therefore, the path $P_T(a, b)$ must go via c . If $d_T(c, a) > \lambda/2$ and $d_T(c, b) > \lambda/2$, then $d_T(a, b) = d_T(a, c) + d_T(c, b) > \lambda$ and since $d_G(a, b) = 1$, we obtain a contradiction with the assumption that the embedding of G in T has distortion λ (condition (2)). Hence $d_T(c, P_G(x, y)) \leq \min\{d_T(c, a), d_T(c, b)\} \leq \lambda/2$, concluding the proof. \square

Lemma 4 *If a graph G λ -embeds into a tree T , then the diameter in G of any cluster C of a layering partition of G is at most 3λ , i.e., $d_G(x, y) \leq 3\lambda$ for any two vertices $x, y \in C$. In particular, $\lambda^*(G, T) \geq D/3$, where D is the maximal diameter of a cluster of a layering partition of G .*

Proof Let $P_G(x, y)$ be a path of G connecting the vertices x and y outside the ball $B_k(s)$, where $k = d_G(s, x) - 1$. Let $P_G(x, s)$ and $P_G(y, s)$ be two shortest paths of G connecting the vertices x, s and y, s , respectively. Let $c \in V(T)$ be the unique vertex in T that is on the intersection $P_T(x, y) \cap P_T(x, s) \cap P_T(y, s)$. Since c belongs to each of the paths $P_T(x, y)$, $P_T(x, s)$, and $P_T(y, s)$, applying Lemma 3 three times, we infer that $d_T(c, P_G(x, y)) \leq \lambda/2$, $d_T(c, P_G(x, s)) \leq \lambda/2$, and $d_T(c, P_G(y, s)) \leq \lambda/2$.

Let a be a closest to c vertex of $P_G(x, s)$ in the tree T , i.e., $d_T(a, c) = d_T(c, P_G(x, s)) \leq \lambda/2$. Let z be a closest to a vertex of $P_G(x, y)$ in T . From condition (1) and previous inequalities, we conclude that $d_G(a, z) \leq d_T(a, z) = d_T(a, P_G(x, y)) \leq d_T(a, c) + d_T(c, P_G(x, y)) \leq \lambda$. Since $z \in P_G(x, y)$ and $P_G(x, y) \cap B_k(s) = \emptyset$, necessarily $d_G(s, z) \geq d_G(s, x) = d_G(s, a) + d_G(a, x)$, yielding $d_G(a, x) \leq d_G(a, z) \leq \lambda$. Analogously, if b is a closest to c vertex of $P_G(y, s)$ in T , then $d_G(b, y) \leq \lambda$ and $d_T(b, c) \leq \lambda/2$. By non-contractibility condition (1) and triangle condition, $d_G(a, b) \leq d_T(a, b) \leq d_T(a, c) + d_T(b, c) \leq \lambda$. Summarizing, we obtain the desired inequality $d_G(x, y) \leq d_G(x, a) + d_G(a, b) + d_G(b, y) \leq 3\lambda$. \square

From Lemma 1 and Corollary 2, we immediately conclude

Corollary 4 *If a graph $G = (V, E)$ λ -embeds into a tree, then there exists an unweighted tree $H = (V, F)$ (without Steiner points) and a $\{0, 1\}$ -weighted tree $H' = (V \cup S', F')$ (with Steiner points), both constructible in linear $O(|V| + |E|)$ time, such that*

$$d_H(x, y) - 2 \leq d_G(x, y) \leq d_H(x, y) + 3\lambda$$

and

$$d_{H'}(x, y) \leq d_G(x, y) \leq d_{H'}(x, y) + 3\lambda$$

for any vertices $x, y \in V$.

This corollary shows that, for any unweighted graph G , it is possible to turn its non-contractive multiplicative low-distortion embedding into a weighted tree to a non-expanding additive low-distortion embedding into a $\{0, 1\}$ -weighted tree. This seems to be an interesting result on its own (note that the additive distortion of embedding general finite metrics into trees can be approximated within a factor of 3 [1, 8]).

Since the largest diameter D of a cluster in \mathcal{LP} can be computed in at most $O(|V||E|)$ time, from Corollaries 1 and 3, we obtain:

Corollary 5 *If a graph $G = (V, E)$ λ -embeds into a tree, then there exists a uniformly weighted tree $H_w = (V, F, w)$ (without Steiner points) and a uniformly weighted tree $H'_w = (V \cup S', F', w)$ (with Steiner points), both constructible in $O(|V||E|)$ time, such that*

$$d_G(u, v) \leq d_{H_w}(u, v) \leq (3\lambda + 1)(d_G(u, v) + 2)$$

and

$$d_G(u, v) \leq d_{H'_w}(u, v) \leq (3\lambda + 1)(d_G(u, v) + 1)$$

for any vertices u, v of G .

Note that, although the topologies H and H' of trees H_w and H'_w can be constructed in linear $O(|V| + |E|)$ time, we need to compute the weights $w = D + 1$ and $w = (D + 1)/2$ assigned to each edge of H and H' , and this requires $O(|V||E|)$ time.

Corollary 5 implies already that there exists a factor 12 approximation algorithm (resp., factor 8 approximation if Steiner points are used) for the problem of non-contractive embedding an unweighted graph into a tree with minimum multiplicative distortion. Below we show that, by strengthening the result of Lemma 4, one can improve the approximation ratio from 12 to 9 and from 8 to 6.

Lemma 5 *Assume that $G = (V, E)$ λ -embed into a tree T , let $C = L_j^i \in \mathcal{LP}$ be a cluster of a layering partition of G and v be an arbitrary vertex of C . Then, for any neighbor $v' \in L^{i-1}$ of v and any vertex $u \in C$, we have $d_G(v', u) \leq \max\{3\lambda - 1, 2\lambda + 1\}$.*

Proof Let $c \in V(T)$ be the nearest common ancestor in the tree T (rooted at s) of all vertices of cluster $C = L_j^i$. Let x and y be two vertices of C separated by c . Let $P_G(x, y)$ be a path of G connecting vertices x and y outside the ball $B_{i-1}(s)$. Then, as in the proof of Lemma 4, we have $d_T(c, P_G(x, y)) \leq \lambda/2$.

Pick an arbitrary vertex $v \in C$ and a shortest path $P_G(v, s)$ connecting v with s in G . Since c separates v from s in T , by Lemma 3, $d_T(c, P_G(v, s)) \leq \lambda/2$ holds. Let a_v be a closest to c vertex of $P_G(v, s)$ in the tree T . Then, $d_T(a_v, P_G(x, y)) \leq d_T(a_v, c) + d_T(c, P_G(x, y)) \leq \lambda$. The choice of the path $P_G(x, y)$ and inequality (1) imply that $d_G(a_v, v) \leq d_G(a_v, P_G(x, y)) \leq d_T(a_v, P_G(x, y)) \leq \lambda$.

Consider now an arbitrary vertex $u \in C$, $u \neq v$. By the triangle inequality and (1), we have $d_G(a_v, a_u) \leq d_T(a_v, a_u) \leq d_T(a_v, c) + d_T(a_u, c) \leq \lambda$ and, therefore, $d_G(a_v, u) \leq d_G(a_v, a_u) + d_G(a_u, u) \leq 2\lambda$.

Let $v' \in L^{i-1}$ be a neighbor of v in $P_G(v, s)$. If $a_v = v$, then $d_G(v, u) = d_G(a_v, u) \leq 2\lambda$, i.e., $d_G(v', u) \leq d_G(v, u) + 1 \leq 2\lambda + 1$. Otherwise, if $a_v \neq v$, then $d_G(v', u) \leq d_G(v', a_v) + d_G(a_v, u) \leq \lambda - 1 + 2\lambda = 3\lambda - 1$, establishing the required inequality $d_G(v', u) \leq \max\{3\lambda - 1, 2\lambda + 1\}$. \square

To make the embedding of G into the tree H non-contractive, it suffices to assign the same length $\ell := \max\{3\lambda - 1, 2\lambda + 1\}$ to each edge of H and get a uniformly weighted tree $H_\ell = (V, F, \ell)$.

Corollary 6 *For any vertices u, v of a graph G which λ -embeds into a tree, we have $d_G(u, v) \leq d_{H_\ell}(u, v) \leq \max\{3\lambda - 1, 2\lambda + 1\}(d_G(u, v) + 2)$.*

The tree H_ℓ provides a 9-approximation to the problem of non-contractive embedding an unweighted graph into a tree with minimum multiplicative distortion. Note that the tree H_ℓ does not have Steiner points. If we allow Steiner points, a better approximation can be achieved. For this, we simply assign the same length $\ell := \frac{3\lambda}{2}$ to each edge of H' and get a uniformly weighted tree H'_ℓ .

Corollary 7 *For any vertices u, v of a graph G which λ -embeds into a tree, we have $d_G(u, v) \leq d_{H'_\ell}(u, v) \leq 3\lambda(d_G(u, v) + 1)$.*

For a given graph $G = (V, E)$, we do not know λ in advance; however, we know from Lemma 4 that $\lambda^*(G, T) \geq D/3$. Therefore, the length ℓ , which needs to be assigned to each edge of the tree H (which is defined in a canonical way, independently of the value of λ), can be found as follows: $\ell = \max\{d_G(u, v) : uv \text{ is an edge of } H\}$. The length ℓ , which needs to be assigned to each edge of tree H' , can be found as follows: $\ell = \frac{1}{2} \max\{D, \max\{d_G(u, v) : uv \text{ is an edge of } H\}\}$. Hence, ℓ can be computed in $O(|V||E|)$ time. Our main result of this section is the following algorithm and theorem.

Algorithm APPROXIMATION BY TREE METRIC

Input: A graph $G = (V, E)$, a root vertex s and the corresponding layering partition $\mathcal{LP} = \{L_1^i, \dots, L_{p_i}^i : i = 0, 1, \dots, r\}$ of G

Output: Trees H, H', H_ℓ , and H'_ℓ for G

1. Set initially $H := H' := (V, \emptyset)$.
2. **For** $i = r$ **downto** 1 **do**
3. **For** each cluster C from $\{L_1^i, \dots, L_{p_i}^i\}$ **do**
4. Pick a vertex x_C in L^{i-1} which has a neighbor in C .
5. Add to H edges $\{vx_C : v \in C\}$.
6. Add to H' a Steiner point p_C and edges $\{vp_C : v \in C\}$ and $p_C x_C$.
7. Set $\ell := \max\{d_G(u, v) : uv \text{ is an edge of } H\}$.
8. Set $H_\ell := H$ and assign length ℓ uniformly to all edges of H_ℓ .
9. Set $\ell := \frac{1}{2} \max\{D, \ell\}$, where D is the largest diameter of a cluster from \mathcal{LP} .
10. Set $H'_\ell := H'$ and assign length ℓ uniformly to all edges of H'_ℓ .
11. Return trees H, H', H_ℓ , and H'_ℓ .

Theorem 1 *There exists a factor 6 approximation algorithm with running time $O(|V||E|)$ for the optimal multiplicative distortion $\lambda^*(G, T)$ of non-contractive embedding an unweighted graph G into a tree.*

Our 6-approximation algorithm improves previously known 100-approximation [4] and 27-approximation [3] algorithms. In fact, the approximation ratio 6 holds only for adjacent vertices of G . It decreases when distances in G increase. For vertices at distance ≥ 2 , the ratio is ≤ 4.5 . For vertices at distance ≥ 3 , the ratio is ≤ 4 . Our tree H_ℓ does not have any Steiner points and the edges of both trees H_ℓ and H'_ℓ are uniformly weighted. The tree H'_ℓ , with Steiner points, is better than the tree H_ℓ only for small graph distances. So, the Steiner points do not really help, confirming A. Gupta's claim [13].

Our technique works also in more general cases. In particular, if an unweighted graph $G = (V, E)$ admits a weighted tree $T = (V \cup S, U)$ with

$$d_G(x, y) \leq \alpha \cdot d_T(x, y) + \beta \quad \text{and} \quad d_T(x, y) \leq \lambda \cdot d_G(x, y) + \delta$$

for all $x, y \in V$, then each cluster of a layering partition of G has diameter at most $3(\alpha(\lambda + \delta) + \beta)$. Moreover, $H = (V, F)$, weighted appropriately, will give a good approximation of T .

4 Minors, Relaxed Minors, and Metric Minors

In this section, we introduce the notions of relaxed minors and α -metric relaxed minors, which, together with layering partitions, are used in the algorithm for approximating the optimal distortion of embedding unweighted graphs into outerplanar graphs (i.e., $K_{2,3}$ -minor free graphs approximation). These concepts and results may be helpful for designing approximation algorithms for embedding graphs into other classes of minor closed graphs.

4.1 Minors and Relaxed Minors

Recall that a graph H is a *minor* of a graph G if a graph isomorphic to H can be obtained from G by contracting some edges, deleting some edges, and deleting some isolated vertices [9]. Notice that the property of being minor is transitive, i.e., if G' is a minor of G and H is a minor of G' , then H is a minor of G . To adapt the concept of minor to our embedding purposes, note that $H = (V', E')$ is a minor of $G = (V, E)$ provided there exists a map $\mu : V' \cup E' \mapsto 2^V$ such that

- (i) For any vertex v of H , $G(\mu(v))$ is connected;
- (ii) For any different vertices v, v' of H , $G(\mu(v)) \cap G(\mu(v')) = \emptyset$;
- (iii) For any edge $e = uv$ of H , $G(\mu(e))$ is a path P_e of G with one end in $G(\mu(u))$ and another end in $G(\mu(v))$;
- (iv) For any vertex v and any edge e of H with $v \notin e$, $P_e \cap G(\mu(v)) = \emptyset$;
- (v') For any two edges $e = (x, y)$, $e' = (u, v)$ of H , the paths P_e and $P_{e'}$ intersect if and only if $\{x, y\} \cap \{u, v\} \neq \emptyset$ and if say, $e = (x, y)$, $e' = (x, w)$ then P_e and $P_{e'}$ intersect only in $\mu(x)$.

Indeed, if such a map μ exists, then contracting each connected subgraph $\mu(v)$, $v \in V'$, to a single vertex v and each path P_e to an edge e , the conditions (ii), (iii), and (v') ensure that the resulting graph will be isomorphic to H , i.e., H is indeed a minor of G . Note that if in (v') two paths P_e and $P_{e'}$ intersect, then they intersect in the subgraph $G(\mu(u))$, where u is the common end of e and e' . In particular, if the edges e, e' are non-incident, then the paths P_e and $P_{e'}$ are disjoint.

For our metric related theory, we will need a weaker notion of minor by allowing intersecting paths to intersect anywhere. A graph $H = (V', E')$ is called a *relaxed minor* of a graph $G = (V, E)$ if there exists a map $\mu : V' \cup E' \mapsto 2^V$ satisfying the conditions (i)–(iv) and the following relaxation of (v'):

- (v) For any two non-incident edges e, e' of H , the paths P_e and $P_{e'}$ are disjoint.

The concept of relaxed minor is weaker than that of minor. For example, the triangle C_3 (3-cycle) is not a minor of any tree, but it is a relaxed minor of the star $K_{1,3}$: μ maps the three vertices of C_3 to the three leaves of $K_{1,3}$ and maps each edge uv of C_3 to the unique path of $K_{1,3}$ between the leaves $\mu(u)$ and $\mu(v)$. The map μ satisfies the conditions (i)–(v), but does not satisfy the condition (v').

Relaxed and α -metric relaxed minors (see Sect. 4.2) are crucial because their existence corresponds to a witness that G cannot be embedded into H -relaxed-minor-free graphs with small distortion (see Proposition 3). Thus it seems important to relate this notion to standard minors. We conjecture that *if the graph H is triangle-free, then the notion of relaxed minor is not weaker than that of minor*. Here we prove a weaker statement. We note that while this leaves this graph-theoretic point not settled, it has no bearing regarding the metric consequences (proving the conjecture will improve the distortion lower bound of Proposition 3 only by a factor of 2). We established a weaker statement which is enough to deal with H of special form: H will be bipartite $H = (V, F; E)$ with every vertex $f \in F$ of degree two. Such *subdivided* graphs H can be seen as a subdivision of an arbitrary graph $H' = (V, E')$ where $(u, v) \in H'$ iff there is a member $f \in F$ such that $(u, f), (v, f) \in E$.

Proposition 2 *If a graph $G = (V, E)$ has a subdivided graph $H = (V', E')$ as a relaxed minor, then G has H as a minor.*

Proof We proceed by induction on the total number of vertices and edges of the graph G . The base case for which $H = G$ is trivial. Let H be a relaxed minor of G and let $\mu : V' \cup E' \mapsto 2^V$ be the map satisfying the conditions (i)–(iv) and (v). Suppose, by way of contradiction, that H is not a minor of G , in particular, that μ does not satisfy the condition (v'). For each edge uv of H , we will denote by $u^* \in \mu(u)$ and $v^* \in \mu(v)$ the end vertices of the path P_e . Note that we can suppose that each such path P_e , $e \in E'$, intersects the connected subgraphs $G(\mu(u))$ and $G(\mu(v))$ in a single vertex. Indeed, if say P_e intersects $G(\mu(u))$ and/or $G(\mu(v))$ in several vertices, then let x be the last vertex of $P_e \cap \mu(u)$ while moving along P_e from u^* to v^* and let y be the first vertex of $P_e \cap \mu(v)$ while moving from x to v^* . Replacing in the definition of μ the path P_e by its subpath P'_e between x and y , we will obtain a map μ which still satisfies the definition of relaxed minors and such that $|P'_e \cap \mu(u)| = 1$ and $|P'_e \cap \mu(v)| = 1$. So, we will further assume that μ obeys this additional condition, i.e., $P_e \cap \mu(u) = \{u^*\}$ and $P_e \cap \mu(v) = \{v^*\}$ for any edge $e = uv$ of H .

We assert now that for each vertex v of the graph H , the subgraph of G induced by $\mu(v)$ consists of a single vertex. If this is not the case, then let G' be the graph obtained from G by contracting the connected subgraph $G(\mu(v))$ to a single vertex v' . Then, G' is a minor of G . Denote by ψ the map from G to G' defining this contraction, i.e., $\psi(u) = u$ if $u \notin \mu(v)$ and $\psi(u) = v'$, otherwise. Then, the composition of μ with ψ is a map from H to G' satisfying the conditions (i)–(v), i.e., H is a relaxed minor of G' . By induction assumption, H is a minor of G' and therefore must be a minor of G as well, contrary to our assumption. Therefore, for each vertex v of H , the set $\mu(v)$ consists of a single vertex of G , which we will further denote by v^* . We can also suppose that the paths P_e , $e \in E'$, are induced paths of G , otherwise we can shortcut them without violating the conditions (i)–(v). Similarly, it is easy to observe that no edge of G is used by more than one path of the form P_e , as otherwise, such an edge can be contracted leaving H as a relaxed minor in the contracted graph.

Since H is a subdivided graph, H is of the form $(V, F; E)$, where each vertex $f \in F$ has degree 2. Since μ does not satisfy the condition (v') for minors, there exist two incident edges $e = uv$ and $e' = uw$ of H such that the paths P_e and $P_{e'}$ of G intersect in other vertices except the vertex $u^* = \mu(u)$. Suppose first that this happens for some $u = f \in F$. Namely, P_e and $P_{e'}$ intersect in u^* and in addition in some $x \in V(G)$, where x is the closest along P_e to u^* . It is easy to see that one can change μ so that to map $\mu(u)$ to x , to map e to the suffix of P_e from x to v^* , and $P_{e'}$ to the suffix from x to w^* , while still having H as a relaxed minor of G . However, now the first edge (u^*, u') in the former path P_e is not used anymore and can be deleted, thus induction ends the proof. We conclude that for every $f \in F$ the two paths that correspond to the edges adjacent to f , intersect only in their ends points $\mu(f)$.

Assume then, that there is vertex in H , $u \in V$ and two edges $e_1 = (u, f_1)$, $e_2 = (u, f_2)$ such that the paths $\mu(e_1) = P_1$ and $\mu(e_2) = P_2$ intersect also at $x \neq u^*$. Let u^*u_1 , u^*u_2 be the first edges in the paths P_1 , P_2 , respectively, and note that $u_1, u_2 \neq x$ otherwise, if say $u_1 = x$, then the edge u^*x could be contracted (and the map μ

changed accordingly), preserving H as a relaxed minor of the smaller resulting graph, implying the result.

Now, the edge (u^*, u_1) cannot be contracted only if u_1 is used by another path $\mu(e_3) = P_3$ for some $e_3 \in H$ that is not adjacent to u (this is allowed as H is a relaxed minor). However, as P_3 intersects P_1 , it must be the case that e_3 is adjacent to e_1 at f_1 . This, however, contradicts our conclusion before that the two paths adjacent to any $\mu(f)$, $f \in F$ intersect only at their end points. This contradiction completes the proof of the proposition. \square

4.2 α -Metric Relaxed Minors

We say that two sets A, B of a graph G are α -far if $\min\{d_G(a, b) : a \in A, b \in B\} > \alpha$. For $\alpha \geq 1$, we call a graph $H = (V', E')$ an α -metric relaxed minor of a graph $G = (V, E)$ if there exists a map $\mu : V' \cup E' \mapsto 2^V$ satisfying the conditions (i)–(v) (i.e., H is a relaxed minor of G) and the following stronger version of condition (v):

(v⁺) For any two non-incident edges $e = uv$ and $e' = u'v'$ of H , the sets $\mu(u) \cup P_e \cup \mu(v)$ and $\mu(u') \cup P_{e'} \cup \mu(v')$ are α -far in G .

Note that while the notion of a relaxed minor is a relaxation of the standard graph-minor (and is weaker), the notion of metric relaxed minor is generally stronger.

To motivate the concept of α -metric relaxed minor, we establish first the following basic property of embeddings with (multiplicative) distortion $\leq \alpha$ of unweighted graphs G into (possibly weighted) graphs G' .

Let φ be an embedding of a graph $G = (V, E)$ into a graph $G' = (V', E')$ having distortion at most α . For a set $S \subseteq V$ inducing a connected subgraph $G(S)$ of G , we denote by $[\varphi(S)]$ a union of shortest paths of G' running between each pair of vertices of $\varphi(S)$ which are images of adjacent vertices of $G(S)$, one shortest path per pair.

Lemma 6 *If a graph G α -embeds into a graph G' and two edges $e_1 = a_1a_2$ and $e_2 = b_1b_2$ are α -far in G , then $[\varphi(e_1)] \cap [\varphi(e_2)] = \emptyset$. More generally, if two sets of vertices A, B induce connected subgraphs of G and are α -far, then $[\varphi(A)] \cap [\varphi(B)] = \emptyset$.*

Proof For a vertex v of G , let $v^* = \varphi(v)$. Suppose, by way of contradiction, that the shortest paths $P_{e_1} = [\varphi(e_1)]$ between a_1^*, a_2^* and $P_{e_2} = [\varphi(e_2)]$ between b_1^*, b_2^* intersect in a vertex x . Since

$$\begin{aligned} 1 &= d_G(a_1, a_2) \leq d_{G'}(a_1^*, a_2^*) \leq \alpha \cdot d_G(a_1, a_2) = \alpha, \\ d_{G'}(a_1^*, a_2^*) &= d_{G'}(a_1^*, x) + d_{G'}(x, a_2^*), \\ 1 &= d_G(b_1, b_2) \leq d_{G'}(b_1^*, b_2^*) \leq \alpha \cdot d_G(b_1, b_2) = \alpha, \\ d_{G'}(b_1^*, b_2^*) &= d_{G'}(b_1^*, x) + d_{G'}(x, b_2^*), \end{aligned}$$

we conclude that

$$\min\{d_{G'}(a_1^*, x), d_{G'}(x, a_2^*)\} \leq \alpha/2 \text{ and } \min\{d_{G'}(b_1^*, x), d_{G'}(x, b_2^*)\} \leq \alpha/2.$$

Suppose, without loss of generality, that $d_{G'}(a_1^*, x) \leq \alpha/2$ and $d_{G'}(b_1^*, x) \leq \alpha/2$. Since $d_G(a_1, b_1) \leq d_{G'}(a_1^*, b_1^*) \leq d_{G'}(a_1^*, x) + d_{G'}(b_1^*, x) \leq \alpha/2 + \alpha/2 \leq \alpha$, we obtain a contradiction with the assumption that the edges $e_1 = a_1a_2$ and $e_2 = b_1b_2$ are α -far in G .

To establish the second assertion, suppose, by way of contradiction, that $[\varphi(A)] \cap [\varphi(B)] \neq \emptyset$. From the definition of the sets $[\varphi(A)]$ and $[\varphi(B)]$, we conclude that $[\varphi(e_1)] \cap [\varphi(e_2)] \neq \emptyset$ for an edge e_1 of $G(A)$ and an edge e_2 of $G(B)$. From the first part of the proof, we know that the edges e_1 and e_2 cannot be α -far, thus the sets A and B cannot be α -far either. \square

We will show now that under some general conditions on H , the presence in a graph G of an α -metric relaxed minor isomorphic to H is an obstacle for embedding G into a H -minor free graph with distortion at most α .

Proposition 3 *If a subdivided 2-connected graph $H = (V', E')$ is an α -metric relaxed minor of a graph $G = (V, E)$, then any embedding of G into an H -minor free graph requires distortion $> \alpha$.*

Proof Suppose, by way of contradiction, that G has an embedding φ with distortion $\leq \alpha$ into an H -minor free graph G' . Let $\mu : V' \cup E' \mapsto 2^V$ be a map showing that H is an α -metric relaxed minor of G . Before deriving a contradiction with this assumption, we consider some properties of maps φ and μ . First, note that we can extend φ from the vertex-set V of G to the edge-set E by associating with each edge e of G the shortest path $P_e := [\varphi(e)]$ of G' . Pick any vertex v of H . Then, $\varphi(\mu(v))$ is a connected subgraph of G' because each of the maps μ and φ maps connected subgraphs to connected subgraphs. Moreover, from Lemma 6, we know that φ maps two α -far connected subgraphs of G to two disjoint subgraphs of G' . As to the map μ , we assert that it satisfies the following two conditions:

- (ii⁺) For any two different vertices v, v' of H , the sets $\mu(v)$ and $\mu(v')$ are α -far;
- (iv⁺) For any vertex v and any edge e of H with $v \notin e$, the sets $\mu(v)$ and $\mu(e) = P_e$ are α -far.

Since H is 2-connected, any two distinct vertices v, v' belong to a common cycle of H . Since H is triangle-free, v and v' belong to two non-incident edges e, e' of this cycle. Applying property (v⁺) to e and e' , we conclude that $\mu(v)$ and $\mu(v')$ are α -far, establishing (ii⁺). Analogously for (iv⁺), if $v \notin e$ then, by 2-connectivity of H , we can find a cycle passing via v and e . Since G is triangle-free, one of two edges of this cycle containing v , say e' , is not incident to e . Again, applying the condition (v⁺) to the edges e and e' , we conclude that the sets $\mu(v)$ and P_e are α -far, establishing (iv⁺).

Now, we define the following map $\nu : V' \cup E' \mapsto 2^{V(G')}$ from H to G' . For each vertex $v \in V'$, we set $\nu(v) = \varphi(\mu(v))$. For each edge $e = uv$ of H , $\mu(e) = P_e$ is a path of the graph G with end-vertices $u^* \in \mu(u)$ and $v^* \in \mu(v)$. Each edge f of P_e is mapped by φ to a path $\varphi(f)$ of G' . Define $\nu(e)$ to be any path of G' between the vertices $u' = \varphi(u^*)$ and $v' = \varphi(v^*)$ contained in the set $\bigcup \{\varphi(f) : f \text{ is an edge of } P_e\}$. From definition of ν and properties of μ and φ , it immediately follows that the map

ν satisfies the conditions (i) and (iii). We will show now that ν satisfies the conditions (ii), (iv), and (v) as well. To verify (ii), pick two distinct vertices u, v of H .

By condition (ii⁺), the sets $\mu(u)$ and $\mu(v)$ are α -far, thus the second assertion of Lemma 6 implies that the sets $\nu(u) = \varphi(\mu(u))$ and $\nu(v) = \varphi(\mu(v))$ are disjoint, thus showing (ii). Analogously, if v is a vertex and e is an edge of H with $v \notin e$, then, by (iv⁺), the sets $\mu(v)$ and $P_e = \mu(e)$ are α -far, thus, by Lemma 6, the sets $\nu(v) = \varphi(\mu(v))$ and $\varphi(P_e)$ are disjoint. Since $\nu(e) \subseteq \varphi(P_e)$, the sets $\nu(v)$ and $\nu(e)$ are disjoint as well, establishing (iv). The last condition (v) can be derived in a similar way by using (v⁺) and Lemma 6. Hence, the map ν satisfies the conditions (i)–(v), thus H is a relaxed minor of G' . Since H is triangle-free, by Proposition 2, H is a minor of G' , contrary to the assumption that the graph G' is H -minor free. This concludes the proof of the proposition. \square

4.3 Lower Bounds for α -Embeddings into $K_{2,r}$ -Minor Free Graphs

We will use the results of previous section to give lower bounds for the multiplicative distortion of embedding an unweighted graph $G = (V, E)$ into $K_{2,r}$ -minor free (possibly weighted) graphs.

Proposition 4 *If for $\alpha > 1$ a cluster C of a layering partition \mathcal{LP} of a graph G contains $r \geq 3$ vertices v_1^*, \dots, v_r^* that are pairwise $(4\alpha + 2)$ -far, then any embedding φ of G into a $K_{2,r}$ -minor free graph has distortion $> \alpha$.*

Proof Suppose that the layering partition \mathcal{LP} of G was defined with respect to the vertex s and let T be a BFS tree rooted at s . Let k denote the distance from s to any vertex of the cluster C . Since C contains $(4\alpha + 2)$ -far vertices v_1^*, \dots, v_r^* , we conclude that $k \geq 2\alpha + 2$. We will define now a mapping μ from $K_{2,r}$ to G which will allow us to conclude that $K_{2,r}$ is an α -metric relaxed minor of G . Then, since $K_{2,r}$ is 2-connected and triangle-free, Proposition 3 will show that any embedding of G into a $K_{2,r}$ -minor free graph has distortion $> \alpha$. Denote by u_1, \dots, u_r, v, w the vertices of $K_{2,r}$, where v and w are the two vertices of degree r . Finally, denote by e_i the edge vu_i and by f_i the edge wu_i , $i = 1, \dots, r$.

Let P_1, \dots, P_r be the paths of the tree T of length $\alpha + 1$ from the vertices v_1^*, \dots, v_r^* , respectively, towards the root s . Denote by u_1^*, \dots, u_r^* the other end vertices of the paths P_1, \dots, P_r . Let R_1, \dots, R_r be the paths of T of length $\alpha + 1$ from the vertices u_1^*, \dots, u_r^* , respectively, towards s . Denote by w_1^*, \dots, w_r^* the other end vertices of the paths R_1, \dots, R_r . Set $\mu(u_i) := u_i^*$, $\mu(e_i) := P_i$ and $\mu(f_i) := R_i$ for $i = 1, \dots, r$. Let $\mu(v)$ be the connected subgraph of G induced by all (or some) paths connecting the vertices v_1^*, \dots, v_r^* outside the ball $B_{k-1}(s)$. Finally, let $\mu(w) := B_{k-2\alpha-2}(s)$ (clearly, w_1^*, \dots, w_r^* belong to $\mu(w)$); for an illustration, see Fig. 2. From the definition of the map μ and of the layering partition \mathcal{LP} , we immediately conclude that μ satisfies the conditions (i) and (iii). We will show now that μ also satisfies the conditions (ii⁺), (iv⁺), and (v⁺). Since $\mu(v) \subseteq \bigcup_{j \geq k} L^j$, $\mu(w) = B_{k-2\alpha-2}(s)$, and the vertices $u_1^* = \mu(u_1), \dots, u_r^* = \mu(u_r)$ all belong to the sphere $L^{k-\alpha-1}$, we conclude that the μ -images of the vertices of $K_{2,r}$ are pairwise α -far in G , whence μ satisfies the condition (ii⁺). Analogously, from the definition of the layering of G

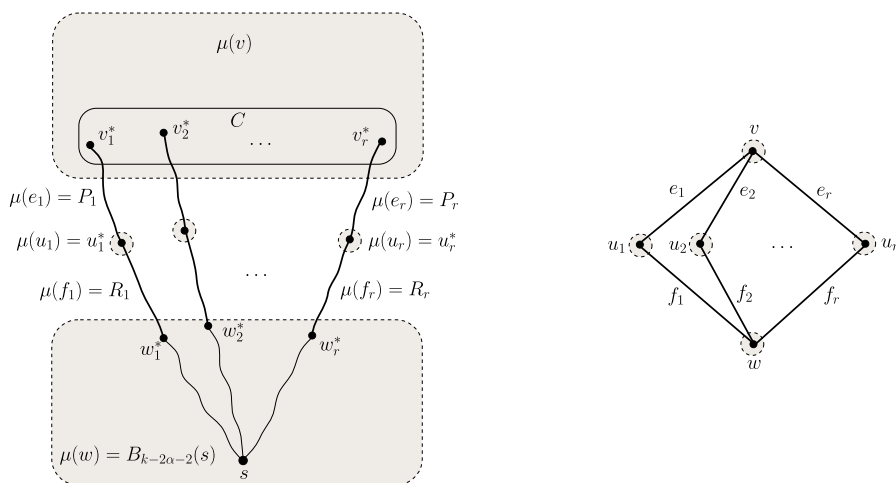


Fig. 2 To the proof of Proposition 4

we conclude that any vertex of $\mu(v)$ is at distance $> \alpha$ from any path $R_i = \mu(f_i)$ and any vertex of $\mu(w)$ is at distance $> \alpha$ from any path $P_i = \mu(e_i)$. If a vertex u_i^* is at distance $\leq \alpha$ from a vertex x of $P_j \cup R_j$ for $j \neq i$, then, by triangle inequality, we obtain $d_G(v_i^*, v_j^*) \leq d_G(v_i^*, u_i^*) + d_G(u_i^*, x) + d_G(x, v_j^*) \leq \alpha + 1 + \alpha + d_G(v_j^*, x)$. Since $x \neq w_j^*$, $d_G(v_j^*, x) \leq 2\alpha + 1$, yielding $d_G(v_i^*, v_j^*) \leq \alpha + 1 + \alpha + 2\alpha + 1 = 4\alpha + 2$, contrary to the assumption that v_i^* and v_j^* are $(4\alpha + 2)$ -far. This contradiction shows that μ satisfies the condition (iv^+) . It remains to show that μ also satisfies the condition (v^+) , namely that any two paths P_i and R_j with $i \neq j$ are α -far. If $d_G(x, y) \leq \alpha$ for $x \in P_i \setminus \{v_i^*, u_i^*\}$ and $y \in R_j \setminus \{u_j^*, w_j^*\}$, then $d_G(v_i^*, v_j^*) \leq d_G(v_i^*, x) + d_G(x, y) + d_G(y, v_j^*) \leq \alpha + \alpha + 2\alpha + 1 \leq 4\alpha + 1$, contrary to the assumption that v_i^* and v_j^* are α -far. This contradiction shows that μ satisfies (v^+) , i.e., indeed $K_{2,r}$ is an α -metric relaxed minor of G . \square

Notice that outerplanar graphs are exactly the graphs which do not contain $K_{2,3}$ and K_4 minors. From Proposition 4, we immediately obtain the following corollary:

Corollary 8 *If for $\alpha > 1$ a cluster C of a layering partition of a graph G contains three vertices that are pairwise $(4\alpha + 2)$ -far, then any embedding φ of G into an outerplanar graph has distortion $> \alpha$.*

5 Approximation Algorithm for Embedding Graph Metrics into Outerplanar Graphs

We present now the algorithm for constant-factor approximation of the distortion of the best embedding of an unweighted graph into outerplanar metrics. Let λ be the best such multiplicative distortion for an input graph G . We first study the structure of a layered partition of G .

5.1 Small, Medium, and Big Clusters

Let $G = (V, E)$ be the input graph and consider a layering partition \mathcal{LP} of G into clusters. We assume that $\lambda \geq 1$ is so that each cluster C of \mathcal{LP} contains at most two vertices which are $(4\lambda + 2)$ -far (otherwise, by Corollary 8, the optimal distortion of embedding G into an outerplanar graph is larger than λ). Set $\Lambda := 4\lambda + 2$. We call a cluster C *bifocal* if the maximum number of pairwise Λ -far vertices of C is exactly two. Let c_1 and c_2 be two Λ -far vertices of such a cluster C . In addition, let $C_1 = \{x \in C : d_G(x, c_1) \leq d_G(x, c_2)\}$ and $C_2 = \{x \in C : d_G(x, c_2) \leq d_G(x, c_1)\}$, and call C_1 and C_2 the *cells* of C centered at c_1 and c_2 , respectively (in what follows, we will suppose that c_1 and c_2 form a diametral pair of C , i.e., $d_G(c_1, c_2) = \text{diam}(C) = \max\{d_G(u, v) : u, v \in C\}$). If $\text{diam}(C) \leq \Lambda$ (i.e., C is not bifocal), then the cluster C is called *small*. Then C has a unique cell centered at an arbitrary vertex of C . A bifocal cluster C is called *big* if $\text{diam}(C) > 16\lambda + 12$, otherwise, if $\Lambda < \text{diam}(C) \leq 16\lambda + 12$, then C is called a *medium* cluster. An *almost big cluster* is a medium cluster C such that $\text{diam}(C) > 16\lambda + 10$. We say that a cluster C is Δ -separated if C is bifocal with cells C_1 and C_2 and $d_G(u, v) > \Delta$ for any $u \in C_1$ and $v \in C_2$. Further, we will set $\Delta := 8\lambda + 6$. We say that a bifocal cluster C' is *spread* if its father C is bifocal and both cells C_1, C_2 of C are adjacent to C' . Finally, we say that two disjoint sets A and B are *adjacent* in G if there exists an edge of G with one end in A and another end in B .

Lemma 7 *If C is a bifocal cluster of a layering partition \mathcal{LP} of G , then the diameter of each of its cells C_1 and C_2 is at most 2Λ .*

Proof Let $x, y \in C_1$. Since C contains no vertex z such that $\{c_1, c_2, z\}$ are pairwise $(4\lambda + 2)$ -far, $d_G(x, c_1) \leq 4\lambda + 2$ and $d_G(y, c_1) \leq 4\lambda + 2$. Therefore, by triangle inequality, $d_G(x, y) \leq d_G(x, c_1) + d_G(c_1, y) \leq 8\lambda + 4 = 2\Lambda$. \square

Lemma 8 *If C is a bifocal cluster of a layering partition \mathcal{LP} of G such that $\text{diam}(C) = d_G(c_1, c_2) > 12\lambda + 6$, then C has the following properties:*

- (i) C is $(\text{diam}(C) - 2\Lambda - 1)$ -separated, in particular $C_1 \cap C_2 = \emptyset$;
- (ii) $\text{diam}(C_1) \leq \Lambda$ and $\text{diam}(C_2) \leq \Lambda$.

If C is a big cluster, then C is $(8\lambda + 8)$ -separated, and if C is an almost big cluster, then C is $(8\lambda + 6)$ -separated. In particular, big and almost big clusters are Δ -separated. If C is a big or an almost big cluster, then $\text{diam}(C_1) \leq \Lambda$ and $\text{diam}(C_2) \leq \Lambda$.

Proof Since the cluster C is bifocal, from the definition of its cells we conclude that, for any two vertices $u \in C_1$ and $v \in C_2$, $d_G(u, c_1) \leq 4\lambda + 2$ and $d_G(v, c_2) \leq 4\lambda + 2$. Therefore, $12\lambda + 6 < \text{diam}(C) = d_G(c_1, c_2) \leq d_G(c_1, u) + d_G(u, v) + d_G(v, c_2) \leq d_G(u, v) + 8\lambda + 4$, showing that $d_G(u, v) > \text{diam}(C) - 2\Lambda - 1$ and $d_G(u, v) > 4\lambda + 2$, whence C is $(\text{diam}(C) - 2\Lambda - 1)$ -separated as well as $(4\lambda + 2)$ -separated. Furthermore, from $d_G(u, v) \geq \text{diam}(C) - 8\lambda - 4$, we obtain that any big cluster (i.e., a cluster C with $\text{diam}(C) > 16\lambda + 12$) is $(8\lambda + 8)$ -separated and any almost big cluster (i.e., a cluster C with $16\lambda + 10 < \text{diam}(C) \leq 16\lambda + 12$) is $(8\lambda + 6)$ -separated. If

C_1 contains two vertices x, y with $d_G(x, y) > 4\lambda + 2$, then the vertices x, y , and c_2 are pairwise $(4\lambda + 2)$ -far, contradicting the assumption that C is bifocal. \square

Given a cluster C located at distance k from the root s and its son C' (in the tree Γ), we call the union of C with the connected component of $G(V \setminus B_k(s))$ containing C' the CC' -fiber of G and denote it by $\mathcal{F}(C, C')$. Note that the son-father relation between clusters that we use here and in what follows is with respect to tree Γ .

Lemma 9 *If a cluster C of a layering partition \mathcal{LP} of G is big, then C has a son C' which is a bifocal spread cluster such that contracting the four cells of C and C' (but preserving the edges between different cells), we will obtain a $2K_2$, an induced matching with two edges.*

Proof Let $C = C_1 \cup C_2$ be the partition of C into cells. Pick $x \in C_1$ and $y \in C_2$ and consider a xy path P in the subgraph of G induced by $V \setminus B_{k-1}(s)$, where k is the distance from the root s to all vertices of C . Since C is big, from Lemma 8(i), we conclude that P cannot entirely lie in C . On the other hand, we can assume, without loss of generality, that C has a son C' such that $P \cap C' \neq \emptyset$ and P is entirely included in the CC' -fiber of G . Therefore, in each of the cells C_1 and C_2 one can pick a vertex which is adjacent to a vertex of C' . Let a_1b_1 and a_2b_2 be two edges of G such that $a_1 \in C_1, a_2 \in C_2$, and $b_1, b_2 \in C'$. Since, by Lemma 8, $d_G(a_1, a_2) > \Delta + 2 = 8\lambda + 8$, we conclude that $d_G(b_1, b_2) \geq 8\lambda + 6 > 4\lambda + 2 = \Lambda$, thus C' is bifocal and the vertices b_1 and b_2 belong to different cells C'_1, C'_2 of C' , say $b_1 \in C'_1$ and $b_2 \in C'_2$. Suppose now that G contains an edge uv with $u \in C_1$ and $v \in C'_2$. Then, by Lemma 8(i), we conclude that $8\lambda + 8 = \Delta + 2 < d_G(u, a_2) \leq 1 + d_G(v, b_2) + 1$, whence $d_G(v, b_2) > 8\lambda + 4 = 2\Lambda$. Since $v, b_2 \in C'_2$, we obtain a contradiction with Lemma 7. Therefore, contracting each of the cells C_1, C_2, C'_1, C'_2 into a vertex, we will indeed obtain a $2K_2$. \square

Lemma 10 *If a cluster C' of a layering partition \mathcal{LP} of G is big or almost big, then its father C is bifocal and the neighbors in C of the centers c'_1 and c'_2 of the cells C'_1 and C'_2 of C' belong to different cells of C . In particular, any big or almost big cluster is spread.*

Proof Let z_1 and z_2 be two neighbors of c'_1 and c'_2 , respectively, in C . If C is not bifocal, then $d_G(z_1, z_2) \leq 4\lambda + 2$, whence $d_G(c'_1, c'_2) \leq 4\lambda + 4 < 16\lambda + 10$, contrary to the assumption that C' is big or almost big. Thus, C is bifocal. If z_1 and z_2 belong to the same cell C_1 of C , then $d_G(z_1, z_2) \leq 2\Lambda$, by Lemma 7, and therefore $d_G(c'_1, c'_2) \leq 2\Lambda + 2 < 16\lambda + 10$, leading to the same contradiction as before. \square

Lemma 11 *If a cluster C of a layering partition \mathcal{LP} of G is big, then no son C' of C has a cell adjacent to both cells of C . In particular, no big cluster C has a small son adjacent to both cells of C .*

Proof Let C_1, C_2 be the cells of C . Suppose, by way of contradiction, that two vertices x', y' from the same cell of C' are adjacent to vertices $x \in C_1$ and $y \in C_2$,

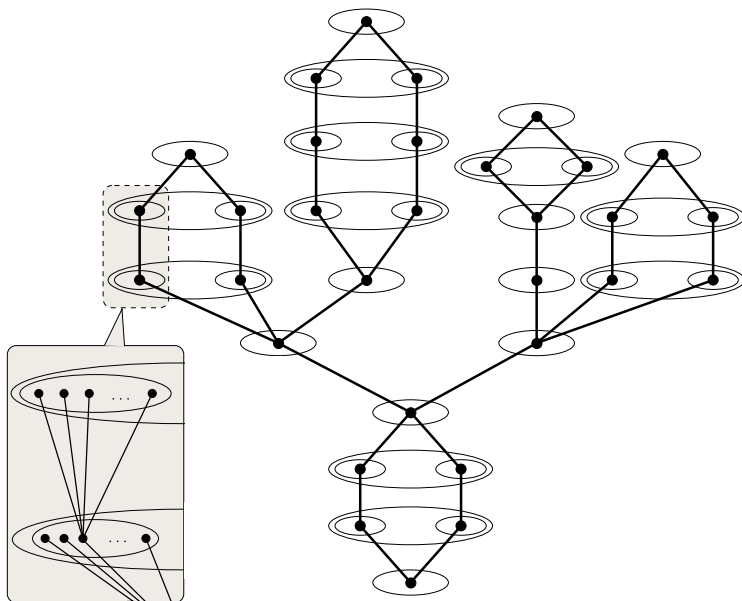


Fig. 3 An outerplanar graph produced by the algorithm APPROXIMATION BY OUTERPLANAR METRIC

respectively. Then, by Lemma 7, $d_G(x, y) \leq 1 + d_G(x', y') + 1 \leq 8\lambda + 4 + 2 = 8\lambda + 6 < \Delta + 2$, contrary to the fact that, according to Lemma 8, the cluster C is $(\Delta + 2)$ -separated. Recall also that a small cluster is not bifocal by definition. \square

5.2 The Algorithm

We continue with the description of an algorithm which, for an input graph G and a current value of “the optimal distortion” λ , either establishes that no embedding with distortion $\leq \lambda$ of G into an outerplanar metric exists or returns such an embedding but with distortion at most $100\lambda + 75$. Namely, given a value of λ such that all clusters of a layering partition \mathcal{LP} of G contain at most two $(4\lambda + 2)$ -far vertices, if some cluster of \mathcal{LP} has two big sons or if this cluster is big and has two spread sons, then any embedding of G in a $K_{2,3}$ -minor free graph requires distortion $> \lambda$, and the algorithm returns the answer “not”. Otherwise, if each cluster has at most one big son and each big cluster has at most one spread son, then the algorithm constructs an outerplanar graph $G' = (V, E')$. Then setting $w := 20\lambda + 15$ as the length of each edge of G' , the inequality $d_G(x, y) \leq d_{G'}(x, y) \leq 5wd_G(x, y)$ holds for any two vertices x, y of V . To construct G' , the algorithm processes the clusters of \mathcal{LP} level by level in increasing order. To ensure that the resulting graph G' is outerplanar and the distortion of the embedding of G into G' is bounded, we need to be precise about how the algorithm “opens” and “closes” the cycles of G' , without allowing cycles to “branch” and without incurring larger and larger distortion. Roughly speaking, small and medium clusters of \mathcal{LP} are used only to open or close cycles of G' or to build

tree-components of G' . Big clusters of \mathcal{LP} are used to build-up the cycles of G' : each cycle C of G' starts and ends with vertices lying in small or medium clusters, all other vertices of C are pairs of centers of cells of big clusters all lying in the same fiber. The remaining vertices of each cell of a big cluster are made adjacent in G' to the neighbor in C of the center of this cell; for an illustration, see Fig. 3. Note that not every outerplanar graph can occur as G' returned by the algorithm because the cycles of G' all have even length and the 2-connected components of G' are edges or cycles. Moreover, no two cycles of G' have a common origin. The precise local rules of constructing G' are provided in lines 3–7 of the algorithm described below.

Algorithm APPROXIMATION BY OUTERPLANAR METRIC

Input: A graph $G = (V, E)$, a layering partition \mathcal{LP} of G , and a value λ .

Output: An outerplanar graph $G' = (V, E')$ or an answer “not”

1. **For** each cluster C of the layering partition \mathcal{LP} **do**
3. **If** C has two big sons or C is big and has two spread sons, **then return** the answer “not”.
3. **Else for** each son C' of C **do**
4. **Case 1:** **If** C' is small, **then** pick in C the center c of a cell of C adjacent to C' and in G' make c adjacent to all vertices of C' .
5. **Case 2:** **If** C' is medium and C is not big, or C' is medium and not spread and C is big, **then** pick in C the center c of a cell of C adjacent to C' and in G' make c adjacent to all vertices of C' .
6. **Case 3:** **If** C' is medium, C is big, and C' is the (unique) spread son of C , **then** in G' make the center c_1 of cell C_1 of C adjacent to all vertices of C' . Additionally, make the center c_2 of cell C_2 of C adjacent to every vertex of C' .
7. **Case 4:** **If** C' is big with cells C'_1, C'_2 , such that C'_1 is adjacent to C_1 and C'_2 is adjacent to C_2 , where C_1 and C_2 are the cells of C with centers c_1 and c_2 , **then** in G' make c_1 adjacent to all vertices of C'_1 and c_2 adjacent to all vertices of C'_2 .

5.3 Correctness of the Algorithm

Now, we formulate the main results establishing the correctness and the approximation ratio of our algorithm. The proofs will be provided in remaining subsections of this section.

Theorem 2 *Let $G = (V, E)$ be an input graph and let $\lambda \geq 1$. If the algorithm APPROXIMATION BY OUTERPLANAR METRIC returns the answer “not”, then any embedding of G into a $K_{2,3}$ -minor free graph requires distortion $> \lambda$. Otherwise, if the algorithm returns the outerplanar graph $G' = (V, E')$, then uniformly assigning to its edges weight $w := 20\lambda + 15$, we obtain an embedding of G to G' such that $d_G(x, y) \leq d_{G'}(x, y) \leq 5wd_G(x, y)$ for any two vertices x, y of V . As a result, we obtain a factor 175 approximation of the optimal distortion of embedding a graph distance into an outerplanar metric.*

The proof of this theorem is subdivided into two propositions. We start with a technical result, essentially showing that in both cases when our algorithm returns the answer “not”, any embedding of G into an outerplanar metric requires distortion $> \lambda$:

Proposition 5 *Let C be a big or an almost big cluster having two sons C' , C'' such that the two cells of C can be connected in both CC' - and CC'' -fibers of G . Then, any embedding of G in a $K_{2,3}$ -minor free graph requires distortion $> \lambda$. These conditions are fulfilled in the following two cases: (i) the cluster C is big and has two spread sons; (ii) C has two big sons C' , C'' .*

In particular, if the algorithm returns the answer “not”, then any embedding of G in a $K_{2,3}$ -minor free graph requires distortion $> \lambda$.

Now suppose that the algorithm returns the graph G' . We continue with the basic property of the graph G' , allowing us to analyze the approximation ratio of the algorithm. First, notice that, by construction, G' is outerplanar. Denote by $d_{G'}(x, y)$ the distance in G' between two vertices x and y , where each edge of G' has length $w := 20\lambda + 15$.

Proposition 6 *For each edge xy of the graph G , the vertices x and y can be connected in the graph G' by a path consisting of at most 5 edges, i.e., $d_{G'}(x, y) \leq 5w$. Conversely, for each edge xy of the graph G' , we have $d_G(x, y) \leq 20\lambda + 15$.*

5.4 Proof of Proposition 6

We start with first assertion. First, suppose that the edge xy of G is horizontal, i.e., $d_G(s, x) = d_G(s, y)$. Let C be the cluster of G containing this edge. Then, either C is not big or C is big and x, y belong to the same cell of C . In both cases, by construction of G' , we deduce that x and y will be adjacent in G' to the same vertex from the father C_0 of C , implying $d_{G'}(x, y) = 2w$. Now suppose that xy is vertical, say $x \in C$, $y \in C'$ and C' is a son of C . Denote by C_0 the father of C . Let z be a vertex of C to which y is adjacent in G' . If C is small, medium, or C is big but x and z belong to the same cell, then in G' the vertices z and x will be adjacent to the same vertex x_{C_0} of the father C_0 of C , yielding $d_{G'}(x, y) \leq 3w$. So, suppose that C is big and the vertices z and x belong to different cells C_1 and C_2 of C , say $z \in C_1$ and $x \in C_2$. By Lemma 11, the cluster C' is not small. According to the algorithm, z is the center of the cell C_1 , i.e., $z = c_1$. Note also that x and the center c_2 of its cell are both adjacent in G' to a vertex $x_{C_0} \in C_0$, whence $d_{G'}(x, c_2) = 2w$. If C' is big and say $y \in C'_1$, then since y is adjacent to z in G' , from the algorithm we conclude that a vertex of C'_1 is adjacent in G to a vertex of C_1 . On the other hand, $y \in C'_1$ is adjacent in G to $x \in C_2$. As a consequence, the cell C'_1 is adjacent in G to both cells C_1 and C_2 of C , which is impossible by Lemma 11. So, the cluster C' must be medium. If C has a big son C'' , then since both cells of C are adjacent in G to the medium son C' , we obtain a contradiction with Proposition 5(i). Hence, C cannot have big sons. Moreover, by Proposition 5, C' is the unique spread son of C . According to the algorithm (see Case 3), the centers $z = c_1$ and c_2 of the cells of C are adjacent in G' to a common vertex u from C' , yielding $d_{G'}(z, c_2) = 2w$. As a result, we obtain a path with at most 5 edges connecting the vertices y and x in G : $(y, z = c_1, u, c_2, x_{C_0}, x)$. This concludes the proof of the first assertion of Proposition 6.

We continue with second assertion. Any edge xy of G' runs between two clusters lying in consecutive layers of G (and G'); let $x \in C$ and $y \in C'$, where C is the

father of C' . In G , the vertex y has a neighbor $x' \in C$. Let $x' \neq x$, otherwise there is nothing to prove. If C is not big, then $d_G(x, x') \leq 16\lambda + 12$, whence $d_G(x, y) \leq 16\lambda + 13$, and we are done. So, suppose that the cluster C is big. If x and x' belong to the same cell of C , then Lemma 7 implies that $d_G(x, x') \leq 2\lambda = 8\lambda + 4$, yielding $d_G(x, y) \leq 8\lambda + 5$. Now, suppose that $x \in C_1$ and $x' \in C_2$. By Lemma 11, C' is a medium or a big cluster. If C' is big and $y \in C'_1$, since x and y are adjacent in G' , according to the algorithm, C'_1 contains a vertex that is adjacent in G to a vertex of C_1 . Since $y \in C'_1$ is adjacent in G to $x' \in C_2$, we obtain a contradiction with Lemma 11. Hence C' is a medium cluster. According to the algorithm, x is the center of the cell C_1 and C_1 contains a vertex z adjacent in G to a vertex $v \in C'$. Since $x, z \in C_1$ implies $d_G(x, z) \leq 4\lambda + 2$ and $y, v \in C'$ implies $d_G(y, v) \leq 16\lambda + 12$, we obtain $d_G(x, y) \leq 20\lambda + 15$.

5.5 Proof of Proposition 5

By Proposition 3, it suffices to show that G contains $K_{2,3}$ as a λ -metric relaxed minor. Indeed, suppose that C is a big or an almost big cluster with cells C_1 and C_2 having two sons C', C'' , such that C_1 and C_2 can be connected by a path in each of the CC' - and CC'' -fibers of G . Let $k = d_G(s, C)$. Denote by P' and P'' the shortest two such paths connecting two vertices of C , one in C_1 and another in C_2 , in $\mathcal{F}(C, C')$ and $\mathcal{F}(C, C'')$, respectively. Denote by $x' \in C_1$ and $y' \in C_2$ the end-vertices of P' and by $x'' \in C_1$ and $y'' \in C_2$ the end-vertices of P'' . Clearly, the choice of P' implies $P' \cap C = \{x', y'\}$ and the choice of P'' implies $P'' \cap C = \{x'', y''\}$. Let w' and w'' be middle vertices of P' and P'' , respectively (if one of these paths has odd length, then it has two middle vertices, and we pick one of them). Let a' and b' be the vertices of P' located at distance $\lambda + 1$ (measured in P') from w' , where a' is located between w' and x' and b' is located between w' and y' . Denote by L' the subpath of P' comprised between a' and w' and by R' the subpath of P' comprised between w' and b' . Analogously, for P'' we can define the vertices a'', b'' and the paths L'' and R'' of length $\lambda + 1$ each. Finally, denote by P'_1 and P'_2 the subpaths of P' comprised between a' and x' and between b' and y' . Analogously, define the subpaths P''_1 and P''_2 of P'' . Pick any shortest path M' in G between the vertices x', x'' and any shortest path M'' between y', y'' . Let F' be a subpath of a shortest path $P(x', s)$ from x' to the root s starting with x' and having length 3λ . Analogously, let F'' be a subpath of a shortest path $P(y'', s)$ from y'' to s starting with y'' and having length 3λ . Let J' and J'' be the subpaths of length $\lambda + 1$ of $P(x', s)$ and $P(y', s)$, which continue F' and F'' , respectively, towards s ; see Fig. 4 for an illustration.

Now we are ready to define a mapping $\mu : V(K_{2,3}) \cup E(K_{2,3}) \mapsto V(G)$ certifying that $K_{2,3}$ is a λ -metric relaxed minor of G . Denote the vertices of $K_{2,3}$ by a, b, c, q', q'' , where the vertices q' and q'' are assumed to be adjacent to each of the vertices a, b, c . We set $\mu(a) := \{w'\}$, $\mu(b) := \{w''\}$, $\mu(q') := P'_1 \cup P''_1 \cup M' \cup F' =: Q'$, $\mu(q'') := P'_2 \cup P''_2 \cup M'' \cup F'' =: Q''$, and $\mu(c) := B_{k'} =: S$, where $k' = k - 4\lambda - 1$. Additionally, for each edge of $K_{2,3}$, we set $\mu(aq') := L'$, $\mu(aq'') := R'$, $\mu(bq') := L''$, $\mu(bq'') := R''$, $\mu(q's) := J'$, $\mu(q''s) := J''$. We will call the paths $L', L'', R', R'', P'_1, P'_2, P''_1, P''_2, F', F'', J', J'', M', M''$, the vertices w', w'' , and the set S the *elements* of the map μ . Notice first that each vertex of $K_{2,3}$ is mapped to

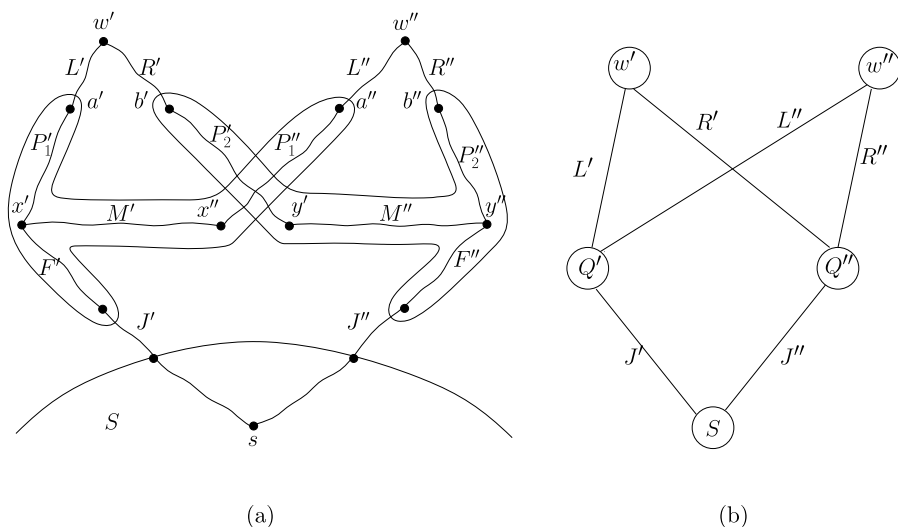


Fig. 4 The elements of the map μ

a connected subgraph of G and each edge of $K_{2,3}$ is mapped to a path of G , thus μ satisfies the conditions (i) and (iii) of a metric relaxed minor. It remains to show that μ satisfies the conditions (ii⁺), (iv⁺), and (v⁺). The proof of this is subdivided into several intermediate results.

Lemma 12 $L' \cup R'$ and $L'' \cup R''$ are shortest paths of G .

Proof Suppose, by way of contradiction, that the vertices a' and b' can be connected in G by a path P_0 shorter than $L' \cup R'$, in particular, $d_G(a', b') \leq 2\lambda + 1$. Since the length of P' is greater than Δ , we conclude that the vertices a' and b' do not belong to the cluster C . From the choice of P' , the path P_0 necessarily contains vertices of $B_{k-1}(s)$, and therefore P_0 necessarily traverses the cluster C . First, suppose that P_0 intersects only one cell of C , say C_1 . Let u be the last vertex of $C_1 \cap P_0$, while moving from a' to b' along P_0 . The length of the subpath Q_0 of P_0 comprised between u and b' is strictly less than the length of P_0 and therefore than the length of $L' \cup R'$. Since b' belongs to the fiber $\mathcal{F}(C, C')$ but does not belong to C , we conclude that necessarily Q_0 is contained in $\mathcal{F}(C, C')$. As a result, the vertices u and y' can be connected in the fiber $\mathcal{F}(C, C')$ by a path $Q_0 \cup P'_2$ shorter than P' , contrary to the choice of P' . Now, suppose that the path P_0 intersects both cells of C . Pick $u \in P_0 \cap C_1$ and $v \in P_0 \cap C_2$. Since u and v can be connected in G by the subpath of P_0 comprised between them, we conclude that $d_G(u, v) \leq 2\lambda + 1$, contrary to the assumption that C is a big or almost big cluster, and thus a Δ -separated cluster. \square

Lemma 13 If $z \in L' \cup R'$, then $d_G(z, C) \geq 4\lambda + 3 - d_G(z, w') \geq 3\lambda + 2$. Analogously, if $z \in L'' \cup R''$, then $d_G(z, C) \geq 4\lambda + 3 - d_G(z, w'') \geq 3\lambda + 2$. In particular, $d_G(w', C) \geq 4\lambda + 3$ and $d_G(w'', C) \geq 4\lambda + 3$.

Proof Note that the length of the subpath of P' between x' and w' is at least $4\lambda + 3$ as C is $(8\lambda + 6)$ -separated. Consequently, the length of the subpath of P' between x' and $z \in L'$ is at least $4\lambda + 3 - d_G(z, w')$. Assume, by way of contradiction, that $d_G(z, C) < 4\lambda + 3 - d_G(z, w')$. First, suppose that $d_G(z, C) = d_G(z, u)$ for $u \in C_1$. Any shortest (z, u) -path $P(u, z)$ lies entirely in the fiber $\mathcal{F}(C, C')$ (and therefore outside the ball $B_{k-1}(s)$). Since $d_G(z, u) < 4\lambda + 3 - d_G(z, w')$, $d_G(z, u)$ is less than the length of the subpath of P' between x' and z . Hence, we conclude that the (u, y') -path consisting of the path $P(u, z)$ followed by the subpath of P' between z and y' is contained in $\mathcal{F}(C, C')$ and is shorter than P' , contrary to the choice of P' . Now suppose that $d_G(z, C) = d_G(z, u)$ for a vertex $u \in C_2$. Since $z \in L'$ and $u \in C$ both belong to $\mathcal{F}(C, C')$, any shortest (z, u) -path $P(u, z)$ also belongs to this fiber. Note that $P(u, z)$ has length $< 4\lambda + 3$ while the subpath of P' between z and y' has length $\geq 4\lambda + 3$. Therefore, the path between u and x' consisting of $P(u, z)$ followed by the subpath of P' between z and x' is shorter than P' and is contained in the fiber $\mathcal{F}(C, C')$, contrary to the choice of P' . \square

Lemma 14 *The set S is λ -far from all elements of μ except J' , J'' and itself.*

Proof From the definition of the layering, it follows that S is λ -far from the paths F' , F'' , and the cluster C . Since any path from a vertex of S to the CC' - and CC'' -fibers traverses the cluster C , we conclude that S is λ -far from the vertices w' , w'' and the paths L' , R' , L'' , R'' , P'_1 , P'_2 , P''_1 , P''_2 . It remains to show that S is λ -far from the paths M' and M'' . Suppose, by way of contradiction, that $d_G(u, v) \leq \lambda$ for $u \in M'$ and $v \in S$. Let $d_G(u, x') \leq d_G(u, x'')$. Since the length of M' is at most $4\lambda + 2$ (by Lemma 8(ii)), we conclude that $d_G(u, x') \leq 2\lambda + 1$, whence $d_G(x', v) \leq 3\lambda + 1$, contrary to the assumption that $d_G(s, x') = k$ and $d_G(s, v) \leq k' = k - 4\lambda - 1$. \square

Lemma 15 *The vertex w' is λ -far from all elements of μ except L' , R' and itself. Analogously, w'' is λ -far from all elements of μ except L'' , R'' and itself.*

Proof From Lemma 13, we conclude that $d_G(w', C) \geq 4\lambda + 3$. Since any path between w' and a vertex of the set $\{w''\} \cup L'' \cup P''_1 \cup R'' \cup P''_2 \cup F'' \cup F' \cup J' \cup J'' \cup S$ traverses the cluster C , we conclude that w' is λ -far from each of these elements of μ . Next we show that w' is λ -far from the paths M' and M'' . Suppose, by way of contradiction, that $d_G(w', u) \leq \lambda$ for a vertex $u \in M'$ and assume, without loss of generality, that u is closer to x' than to x'' , yielding $d_G(x', u) \leq 2\lambda + 1$. But then $d_G(w', x') \leq 3\lambda + 1 < 4\lambda + 3$, contrary to the assumption that $d_G(w', C) \geq 4\lambda + 3$. Finally, we will show that w' is λ -far from the paths P'_1 and P'_2 . Suppose, by way of contradiction, that $d_G(w', u) \leq \lambda$ for a vertex $u \in P'_1$. Let $P(u, w')$ be a shortest (u, w') -path. Obviously, $P(u, w')$ belongs to the CC' -fiber $\mathcal{F}(C, C')$. Hence, replacing in P' the subpath comprised between u and w' (and comprising L') by $P(u, w')$, we obtain a shorter path connecting x' and y' in $\mathcal{F}(C, C')$. This contradiction shows that w' is λ -far from P'_1 and P'_2 . \square

Lemma 16 *Each of the paths L' and R' is λ -far from each of the elements L'' , R'' , P''_1 , P''_2 , J' , J'' of μ . Analogously, L'' and R'' are λ -far from P'_1 , P'_2 , J' , J'' . In particular, the μ -images of any two non-incident edges of $K_{2,3}$ are λ -far.*

Proof The cluster C separates the CC' -fiber containing $L' \cup R'$ from the rest of the graph. Therefore, any path connecting a vertex $u \in L' \cup R'$ to a vertex $v \in L'' \cup R'' \cup P_1'' \cup P_2'' \cup J' \cup J''$ traverses C . Since $d_G(u, C) \geq 4\lambda + 3 - d_G(u, w') \geq 4\lambda + 3 - \lambda - 1 = 3\lambda + 2 > \lambda$, we conclude that $d_G(u, v) > \lambda$. \square

Lemma 17 *The set Q' is λ -far from the paths R' , R'' and J'' . Analogously, the set Q'' is λ -far from the paths L' , L'' , and J' .*

Proof That J'' is λ -far from P_1' and P_1'' follows from the definition of J'' and the fact that the cluster C separates J'' from $P_1' \subseteq \mathcal{F}(C, C')$ and $P_1'' \subseteq \mathcal{F}(C, C'')$. Now suppose that $d_G(u, v) \leq \lambda$ for $u \in M' \cup F'$ and $v \in J'' \setminus S$. If $u \in F'$, then $d_G(x', u) \leq 3\lambda$, $d_G(v, y'') \leq 4\lambda$ and, by triangle inequality, we conclude that $d_G(x', y'') \leq 8\lambda < \Delta$, contrary to the fact that $x' \in C_1$, $y'' \in C_2$ and the cluster C is Δ -separated. If $u \in M'$, then since M' has length at most $4\lambda + 2$, we conclude that one of the vertices x', x'' , say x' has distance at most $2\lambda + 1$ to u . Then, by triangle inequality, again we conclude that $d_G(x', y'') \leq 7\lambda + 1$, contrary with Δ -separability of C . This shows that Q' and J'' are λ -far.

It remains to show that Q' and $R' \cup R''$ are λ -far. Pick $u \in Q'$ and $v \in R''$. By Lemma 13, any vertex $v \in R''$ is located at distance $\geq 3\lambda + 2$ from the cluster C . Since C separates R'' from P_1' and F' , we conclude that $d_G(u, v) \geq 3\lambda + 2$ for any vertex $u \in P_1' \cup F'$. If $u \in M'$ and $d_G(x', u) \leq d_G(x'', u)$, then $d_G(x', u) \leq 2\lambda + 1$, yielding $d_G(x', v) \leq d_G(x', u) + d_G(u, v) \leq 2\lambda + 1 + d_G(u, v)$. Hence, if $d_G(u, v) \leq \lambda$, we get $d_G(x', v) \leq 3\lambda + 1$, contrary to the fact that $d_G(v, C) \geq 3\lambda + 2$. Finally, suppose that $u \in P_1''$ and $d_G(u, v) \leq \lambda$. Let P_0 be any shortest path between u and v . Replacing the subpath $P''(u, v)$ of P'' comprised between u and v by P_0 , we will obtain a path P shorter than P'' (because $L'' \subset P''(u, v)$ and, by Lemma 12, L'' is a shortest path of length $\lambda + 1$ of G). The path P is completely contained in the union of the fiber $\mathcal{F}(C, C'')$ and the ball $B_k(s)$. Moreover, each time P moves from $\mathcal{F}(C, C'')$ to $B_k(s)$, it traverses the cluster C . Therefore, taking any subpath of P between two vertices from different cells of C and completely contained in $\mathcal{F}(C, C'')$, we will obtain a contradiction with the minimality choice of the path P'' . This contradiction concludes the proof that Q' is λ -far from R' , R'' and J'' . \square

Lemma 18 *The sets Q' and Q'' are λ -far.*

Proof First, notice that M' and M'' are λ -far. Indeed, pick $u \in M'$, $v \in M''$, and suppose, without loss of generality, that $d_G(x', u) \leq d_G(x'', u)$ and $d_G(y'', v) \leq d_G(y', v)$. If $d_G(u, v) \leq \lambda$, then, by triangle inequality, $d_G(x', y'') \leq d_G(x', u) + d_G(u, v) + d_G(v, y'') \leq 2\lambda + 1 + \lambda + 2\lambda + 1 = 5\lambda + 2 < \Delta$, contrary to assumption that C is Δ -separated. In a similar way, one can show that F' and M'' as well as M' and F'' are λ -far: if $d_G(u, v) \leq \lambda$ for $u \in F'$ and $v \in M''$ with $d_G(y'', v) \leq 2\lambda + 1$, then $d_G(x', y'') \leq 3\lambda + \lambda + 2\lambda + 1 = 6\lambda + 1 < \Delta$. Analogously, if F' and F'' are not λ -separated, then $d_G(x', y'') \leq 3\lambda + \lambda + 3\lambda = 7\lambda < \Delta$, a contradiction.

Suppose now that $u \in Q'$, $v \in P_2''$, and $d_G(u, v) \leq \lambda$. Let P_0 be a shortest path of G between u and v . Since C is Δ -separated, P_0 cannot intersect both cells C_1 and C_2 of C . On the other hand, since $v \in P_2'' \subset \mathcal{F}(C, C'')$, the path P_0 necessarily contains

a vertex $v_0 \in C$ such that the whole subpath of P_0 between v_0 and v is contained in $\mathcal{F}(C, C'')$. If $v_0 \in C_1$, then the path constituted by the subpath of P_0 between v_0 and v , followed by the subpath of P_2'' between v_0 and y'' , is completely contained in the fiber $\mathcal{F}(C, C'')$ and is shorter than P'' (because $L'' \cup R''$ has length $2\lambda + 2$), contrary to the minimality choice of P'' . Therefore, necessarily $v_0 \in C_2$, showing also that $P_0 \cap C = P_0 \cap C_2$. Let also u_0 be the first intersection of P_0 with C_2 while moving from u to v .

If $u \in F'$ then $d(x', u) \leq 3\lambda$ and we conclude $d_G(x', u_0) \leq d_G(x', u) + d_G(u, u_0) \leq 3\lambda + \lambda = 4\lambda < \Delta$, contrary to the fact that C is Δ -separated. Analogously, if $u \in M'$ and $d_G(x', u) \leq d_G(x'', u)$, then $d_G(x', u_0) \leq d_G(x', u) + d_G(u, u_0) \leq 2\lambda + 1 + \lambda < \Delta$. If $u \in P_1'$, then the subpath of P' between x' and u followed by the subpath of P_0 between u and u_0 forms a path contained in the fiber $\mathcal{F}(C, C')$ and that is shorter than P' (because $L' \cup R'$ has length $2\lambda + 2$), contrary to the choice of P' . Finally, suppose that $u \in P_1''$. Then, the subpath of P'' between x'' and u followed by the subpath of P_0 between u and u_0 constitute a path contained in the fiber $\mathcal{F}(C, C'')$ and is shorter than P'' , contrary to the choice of P'' . This contradiction shows that the sets Q' and Q'' are λ -far. \square

This establishes the first assertion of Proposition 5. To prove the second assertion of Proposition 5, first suppose that the cluster C is big and C has a big and a medium sons C', C'' such that both cells C_1 and C_2 are adjacent to C'' or that C has two medium sons C', C'' adjacent to both cells of C . By definition of the layering, each vertex of $C' \cup C''$ is adjacent to a vertex of C . If all vertices of C' are adjacent to vertices from the same cell of C , say C_1 , then for any $x', y' \in C'$ we have $d_G(x', y') \leq 2 + 4\lambda + 2$, contrary to the assumption that C' is big. Hence, both cells of C are adjacent to C' , say $x \in C_1$ is adjacent to $x' \in C'$ and $y \in C_2$ is adjacent to $y' \in C'$. By Lemma 11, x' and y' belong to different cells of C' , say $x' \in C_1'$ and $y' \in C_2'$. Let $k := d_G(s, C)$. Since $x', y' \in C'$, the vertices x' and y' are adjacent in $G(V \setminus B_k(s))$ by a path $P(x', y')$. Then $P(x, y) := xx' \cup P(x', y') \cup y'y$ is a path between x and y in the CC' -fiber $\mathcal{F}(C, C')$. Analogously, since both cells C_1 and C_2 are adjacent to C'' , we conclude that two vertices from different cells of C can be connected by a path belonging to the CC'' -fiber, showing that the conditions of Proposition 5 are fulfilled. This concludes the proof in case (i).

Now suppose that C has two big sons C' and C'' . Then C is either a big or an almost big cluster. By Lemma 9, each of the clusters C', C'' is $(8\lambda + 8)$ -separated while the cluster C is $(8\lambda + 6)$ -separated and that its cells C_1 and C_2 have diameters at most Δ . As in previous cases, one can deduce that C_1 is adjacent to one cell of each of the clusters C' and C'' , while C_2 is adjacent to the second cell of these clusters, establishing the case (ii) and concluding the proof of Proposition 5.

5.6 Proof of Theorem 2

The algorithm returns the answer “not” when a cluster C has two big sons or a big cluster C has two spread sons. In this case, by Proposition 5, any embedding of G into a $K_{2,3}$ -minor free graph requires distortion $> \lambda$, whence $\lambda^*(G, \mathcal{O}) > \lambda$. Now suppose that the algorithm returns the outerplanar graph G' weighted uniformly with

$w = 20\lambda + 15$. Notice that in Case 4 of the algorithm, the required matching between the four cells of the big clusters C and C' exists by Lemma 9 and because C' is the unique spread son of C . By Proposition 6, we have $d_G(x, y) \leq 20\lambda + 15 = d_{G'}(x, y)$ for each edge xy of the graph G' . By Lemma 2, we conclude that $d_G(x, y) \leq d_{G'}(x, y)$ for any pair $x, y \in V$. By Proposition 6, for any edge xy of G , the vertices x and y can be connected in G' by a path with at most 5 edges, i.e., $d_{G'}(x, y) \leq 5w = 100\lambda + 75$. By Lemma 1, we conclude that $d_{G'}(x, y) \leq (100\lambda + 75)d_G(x, y)$ for any pair x, y of V . Hence $d_G \leq d_{G'} \leq (100\lambda + 75)d_G$, concluding the proof of Theorem 2.

6 Final Remarks

All our non-algorithmic results hold for infinite graphs as well. Namely, given a connected graph G with an arbitrary number of vertices, we can define a layering partition \mathcal{LP} of G . Then, as in Sects. 2 and 3, the largest diameter of a cluster of \mathcal{LP} can be used to upper bound the optimal (additive or multiplicative) distortion of embedding G into a tree metric. The trees H, H', H_ℓ , and H'_ℓ can be defined as in the algorithm APPROXIMATION BY TREE METRIC and these trees have the same approximation qualities as the analogous trees defined in the finite case (see Corollaries 4–7). Therefore, an infinite graph G admits an embedding into a tree-metric with a finite distortion if and only if the diameters of clusters of an arbitrary layering partition \mathcal{LP} of G are uniformly bounded. Similar conclusions hold for approximation by outerplanar graphs: Propositions 4, 5, and 6 establish in what cases the optimal distortion is $> \lambda$; otherwise, the construction provided by the algorithm APPROXIMATION BY OUTERPLANAR GRAPH returns an infinite outerplanar graph G' into which G embeds with distortion $\leq 100\lambda + 75$. The proof of Proposition 2 (done by induction on the number of vertices and edges of G) seems to be an obstacle to this conclusion. However, if a graph G has a *finite* graph H as a relaxed minor, then one can easily find a finite subgraph G_0 of G which still has H as a relaxed minor, and therefore we can apply Proposition 2 to G_0 instead of G to conclude that G_0 (and therefore G) has H as a minor.

We conclude with two open questions. Our Proposition 4 presents a strong necessary condition for embedding a graph G into a $K_{2,r}$ -minor free metric with a distortion $\leq \lambda$. However, we were not able to provide all structural conditions and to design a constant factor approximation algorithm for this problem. An even more challenging problem is designing a constant factor approximation algorithm for optimal distortion of embedding a graph metric into a K_4 -minor free metric (series-parallel metric).

Note. After this manuscript was submitted to the journal, we learned about the paper by B. Krön and R.G. Möller “Quasi-isometries between graphs and trees,” J. Combin. Theory Ser. B 98 (2008), 994–1013, which using different terminology also establishes a relationship between the diameters of clusters of a layering partition and the parameters of a quasi-isometric embedding of a graph metric into a tree. In particular, facts similar to our Proposition 1 and Lemma 4 were established.

Acknowledgement We would like to acknowledge the two anonymous referees for a careful reading of the first version of the manuscript and helpful suggestions. This research was supported in part by the ANR Blanc ANR-10-BLAN 0116, acronym GGAA and by the Israel Science Foundation grant number 862/10. We are also grateful to Aix-Marseille Université which made this research possible.

References

1. Agarwala, R., Bafna, V., Farach, M., Narayanan, B., Paterson, M., Thorup, M.: On the approximability of numerical taxonomy (fitting distances by tree metrics). *SIAM J. Comput.* **28**, 1073–1085 (1999)
2. Bădoiu, M., Chuzhoy, J., Indyk, P., Sidiropoulos, A.: Low-distortion embeddings of general metrics into the line. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005) Baltimore, MD, USA, May 22–24, pp. 225–233. ACM, New York (2005)
3. Bădoiu, M., Demaine, E.D., Hajiaghayi, M.T., Sidiropoulos, A., Zadimoghaddam, M.: Ordinal embedding: approximation algorithms and dimensionality reduction. In: Proceedings of the 11th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2008), Boston, MA, USA, August 25–27. Lecture Notes in Computer Science, vol. 5171, pp. 21–34. Springer, Berlin (2008)
4. Bădoiu, M., Indyk, P., Sidiropoulos, A.: Approximation algorithms for embedding general metrics into trees. In: Proceedings of the Eighteenth Annual ACM–SIAM Symposium on Discrete Algorithms (SODA 2007), New Orleans, Louisiana, USA, January 7–9, pp. 512–521. ACM/SIAM, New York (2007)
5. Brandstädt, A., Chepoi, V., Dragan, F.: Distance approximating trees for chordal and dually chordal graphs. *J. Algorithms* **30**, 166–184 (1999)
6. Chepoi, V., Dragan, F.: A note on distance approximating trees in graphs. *Eur. J. Comb.* **21**, 761–766 (2000)
7. Chepoi, V., Dragan, F.F., Estellon, B., Habib, M., Vaxès, Y.: Diameters, centers, and approximating trees of δ -hyperbolic geodesic spaces and graphs. In: Proceedings of the 24th Annual ACM Symposium on Computational Geometry (SoCG 2008), June 9–11, College Park, Maryland, USA, pp. 59–68. ACM, New York (2008)
8. Chepoi, V., Fichet, B.: l_∞ -Approximation via subdominants. *J. Math. Psychol.* **44**, 600–616 (2000)
9. Diestel, R.: Graph Theory. Springer, Berlin (2005)
10. Dourisboure, Y., Dragan, F.F., Gavaille, C., Yan, C.: Spanners for bounded tree-length graphs. *Theor. Comput. Sci.* **383**, 34–44 (2007)
11. Dourisboure, Y., Gavaille, C.: Tree-decompositions with bags of small diameter. *Discrete Math.* **307**, 208–229 (2007)
12. Emek, Y., Peleg, D.: Approximating minimum max-stretch spanning trees on unweighted graphs. *SIAM J. Comput.* **38**, 1761–1781 (2008)
13. Gupta, A.: Steiner points in tree metrics don’t (really) help. In: Proceedings of the Twelfth Annual Symposium on Discrete Algorithms (SODA 2001), January 7–9, Washington, DC, USA, pp. 220–227. ACM/SIAM, New York (2001)
14. Indyk, P., Matousek, J.: Low-distortion embeddings of finite metric spaces. In: Goodman, J.E., O’Rourke, J. (eds.) *Handbook of Discrete and Computational Geometry*, 2nd edn., pp. 177–196. CRC Press, Boca Raton (2004)
15. Krauthgamer, R., Lee, J.R.: Algorithms on negatively curved spaces. In: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21–24 October 2006, Berkeley, California, USA, pp. 119–132. IEEE, New York (2006)
16. Rabinovich, Y., Raz, R.: Lower bounds on the distortion of embedding finite metric spaces in graphs. *Discrete Comput. Geom.* **19**, 79–94 (1998)
17. Shavitt, Y., Tankel, T.: Hyperbolic embedding of Internet graph for distance estimation and overlay construction. *IEEE/ACM Trans. Netw.* **16**, 25–36 (2008)
18. Semple, C., Steel, M.: Phylogenetics. Oxford University Press, London (2003)