Constant approximation algorithms for embedding graph metrics into trees and outerplanar graphs[*]

V. Chepoi[1], F. F. Dragan[2], I. Newman[3], Y. Rabinovich[3], and Y. Vaxès[1]

[1] LIF, Universitée d'Aix-Marseille,
13288 Marseille, France, {chepoi,vaxes}@lif.univ-mrs.fr
[2] Computer Science Department, Kent State University,
Kent, OH 44242, USA dragan@cs.kent.edu
[3] Department of Computer Science, University of Haifa,
Mount Carmel, Haifa 31905, Israel {ilan,yuri}@cs.haifa.ac.il

Abstract. We present a simple factor 6 algorithm for approximating the optimal multiplicative distortion of embedding (unweighted) graph metrics into tree metrics (thus improving and simplifying the factor 100 and 27 algorithms of Bădoiu et al. (2007) and Bădoiu et al. (2008)). We also present a constant factor algorithm for approximating the optimal distortion of embedding graph metrics into outerplanar metrics. For this, we introduce a notion of 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 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.

1 Introduction

1.1. Avant-propos. The structure of the shortest-path metrics of special classes of graphs, in particular, graphs families defined by 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 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., [13] for a survey of algorithmic applications of embeddings. 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 (essentially outerplanar metrics).

The study of tree metrics can be traced back to the beginning of the 20th 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 [16], it was observed that certain Internet originated metrics display tree-like properties. It is well known [17] that tree metrics have a simple structure: d is a tree metric iff 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 algorithmical aspects of finding a tree metric of least additive distortion has culminated in the paper [1] (see also [8]), where a 6-approximation algorithm was established (in the notation of [1], their algorithm is a 3 approximation, however, in our more restrictive definition, this 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 study of algorithmic and other aspects of such geometries, see e.g. [7, 14]. The situation with the *multiplicative* distortion is less satisfactory. The best result for general 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 line metrics, it is conceivable that any constant factor approximation for the general metric 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 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 (such an obstacle was used also in [11], and, essentially, in [15]).

1.2. Our results. In this paper, we simplify and improve the construction of [4], using a decomposition procedure developed earlier in [5, 6]. The improved constant is 6 and the running time of the algorithm is linear once the distance matrix is computed. We also introduce the notion of 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 free graph is at least α. This generalizes the obstacle of [4]. Using this newly defined H-obstacle, we show that it is an essential obstacle not only for trees, but also for graphs without $H = K_{2,3}$ minors. We further develop an efficient algorithm which either embeds the input metric induced by a unit-weighted 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 α-metric relaxed minor corresponding to the four-point condition used for embedding into tree-metrics).

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 $x, y \in X$. 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 $x, y \in X$. 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 $x, y \in X$ and, in a similar way, we can define the minimum additive distortion. In this paper, we consider unweighted graphs as input metric spaces and *tree metrics* (trees) or *outerplanar metrics* 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 u and v. A finite metric space (X, d) is called a *tree metric* [17] if it isometrically embeds into a tree, i.e., there exists an 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. Denote by $\mathcal{T}$ the class of tree metric spaces and by $\mathcal{O}$ the class of outerplanar metric spaces.

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. The layering partitions have been introduced in [5, 6] and recently used in a slightly more general forms in both approximation algorithms of [3, 4] and in other similar contexts [7, 9, 10]. Let $G = (V, E)$ be a 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* $L^i = \{u \in V : d(s, u) = i\}$, $i = 0, 1, 2, \ldots, r$. A *layering partition* $\mathcal{L}P(s) = \{L_1^i, \ldots, L_{p_i}^i : i = 0, 1, 2, \ldots, r\}$ of G is a partition of each L^i into *clusters* $L_1^i, \ldots, L_{p_i}^i$ such that two vertices $u, v \in L^i$ belong to the same cluster L_j^i iff they can be connected by a path outside the ball $B_{i-1}(s)$ of radius $i - 1$ centered at s. Let Γ be a graph whose vertex set is the set of all clusters L_j^i in a layering partition $\mathcal{L}P$ and $C = L_j^i$ and $C' = L_{j'}^{i'}$ are adjacent in Γ iff there exist $u \in L_j^i$ and $v \in L_{j'}^{i'}$ such that u and v are adjacent in G. Γ is a tree [6], called the *layering tree* of G. $\mathcal{L}P$ and Γ are computable in linear time [6]. We can construct a new tree $H = (V, F)$ (closely reproducing the global structure of Γ) by identifying for each cluster $C = L_j^i$ an arbitrary vertex $x_C \in L^{i-1}$ (the *support vertex* for cluster C) which has a neighbor in C and by making x_C adjacent in H with all vertices $v \in C$. In what follows, we assume that Γ and H are rooted at s. Let D be the largest diameter of a cluster in $\mathcal{L}P$, i.e., $D := \max_{C \in \mathcal{L}P} \max_{v, u \in C}\{d_G(u, v)\}$. The following result (also implicitly used in [5–7]) shows that the additive distortion of the embedding of G into H is essentially D:

Proposition 1. *If $x, y \in V$, then $d_H(x,y) - 2 \leq d_G(x,y) \leq d_H(x,y) + D$.*

Proof. Let C_x and C_y be the clusters containing x and y. Let C be the nearest common ancestor of C_x and C_y in Γ. For $C \neq C_x$, let $x', y' \in C$ be the ancestors of x and y 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 to $d_\Gamma(C_x, C) + d_\Gamma(C_y, C)$ or to $d_\Gamma(C_x, C) + d_\Gamma(C_y, C) + 2$. By the triangle inequality, $d_G(x,y) \leq d_G(x, x') + d_G(x', y') + d_G(y, y') \leq d_\Gamma(C_x, C) + d_\Gamma(C_y, C) + D \leq d_H(x,y) + D$. By definition of clusters, $d_G(x,y) \geq d_G(x,x') + d_G(y,y') \geq d_H(x,y) - 2$. $\square$

The BFS-tree H preserves the 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. Assigning length $w := D + 1$ to each edge of H, we will get a tree $H_w = (V, F, w)$ in which G embeds with distortion essentially equal to $D + 1 : d_G(u,v) \leq d_{H_w}(u,v) \leq (D+1)(d_G(u,v) + 2) \ \forall u, v \in V$. Adding Steiner points and using edge lengths 0 and 1, H can be 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: $d_{H'}(u,v) \leq d_G(u,v) \leq d_{H'}(u,v) + D \ \forall u, v \in V$. Finally, by replacing each edge in H' with edge of length $w := \frac{D+1}{2}$, we obtain a tree H'_w so that $d_G(u,v) \leq d_{H'_w}(u,v) \leq (D+1)(d_G(u,v) + 1) \ \forall u, v \in V$.

2.2. Embeddings with distortion λ of graph metrics. We continue with two auxiliary standard results about embeddings.

Lemma 1. *If $G = (V, E), G' = (V', E')$ are two graphs, one unweighted and second weighted, and $\varphi : V \mapsto V'$ is a map so that $d_{G'}(\varphi(u), \varphi(v)) \leq \lambda \ \forall uv \in E$, then $d_{G'}(\varphi(x), \varphi(y)) \leq \lambda d_G(x,y) \ \forall x, y \in V$.*

Lemma 2. *If $G = (V, E), G' = (V', E')$ are two graphs, one unweighted and second weighted, and $\varphi : V \mapsto V'$ is a map so that $d_{G'}(\varphi(u), \varphi(v)) \geq d_G(u,v) \ \forall \varphi(u)\varphi(v) \in E'$, then $d_{G'}(\varphi(x), \varphi(y)) \geq d_G(x,y) \ \forall x, y \in V$.*

3 Embedding into trees

We describe now a simple factor 6 algorithm for approximating the optimal distortion $\lambda^* = \lambda^*(G, \mathcal{T})$ 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 (1) $d_G(x,y) \leq d_T(x,y)$ (non-contractibility) and (2) $d_T(x,y) \leq \lambda \cdot d_G(x,y)$ (bounded expansion) 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 $9\lambda^*$. 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 G λ-embeds into a tree, then for any $x, y \in V$, any (x,y)-path $P_G(x,y)$ of G and any vertex $c \in P_T(x,y)$, $d_T(c, P_G(x,y)) \leq \lambda/2$.*

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 G λ-embeds into a tree T, then the diameter in G of any cluster C of a layering partition of G is $\leq 3\lambda$, i.e., $d_G(x,y) \leq 3\lambda$ for any $x, y \in C$. In particular, $\lambda^*(G, \mathcal{T}) \geq D/3$.*

Proof. Let $P_G(x,y)$ be a (x,y)-path of G 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 x, s and y, s, respectively. Let $c \in V(T)$ be the unique vertex of T in $P_T(x,y) \cap P_T(x,s), \cap P_T(y,s)$. 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 $d_T(c, P_G(x,y)), d_T(c, P_G(x,s))$, and $d_T(c, P_G(y,s))$ are $\leq \lambda/2$. 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 (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 $d_G(s,z) \geq d_G(s,y) = d_G(s,a) + d_G(a,x)$, 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 $d_G(x,y) \leq d_G(x,a) + d_G(a,b) + d_G(b,y) \leq 3\lambda$. $\square$

Lemma 1 and the properties of H' imply that one can construct in linear time an unweighted tree $H = (V, F)$ (without Steiner points) and a $\{0,1\}$-weighted tree $H' = (V \cup S', F')$ (with Steiner points), so that $d_H(x,y) - 2 \leq d_G(x,y) \leq d_H(x,y) + 3\lambda$ and $d_{H'}(x,y) \leq d_G(x,y) \leq d_{H'}(x,y) + 3\lambda \; \forall x, y \in V$. Hence, for any graph G, it is possible to turn its non-contractive multiplicative distortion embedding into a weighted tree to a non-expanding additive distortion embedding into a $\{0,1\}$-weighted tree. From properties of the trees H_w and H'_w, we obtain:

Corollary 1. *If $G = (V, E)$ λ-embeds into a tree, then there exists uniformly weighted trees $H_w = (V, F, w)$ and $H'_w = (V \cup S', F', w)$ (without and 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 $d_G(u,v) \leq d_{H'_w}(u,v) \leq (3\lambda + 1)(d_G(u,v) + 1) \; \forall u, v \in V$.*

Corollary 1 implies already that there exists a factor 12 (factor 8 if Steiner points are used) approximation algorithm for considered problem. We will show now that, by strengthening Lemma 4, one can improve the approximation ratio from 12 to 9 and from 8 to 6.

Lemma 5. *If $G = (V, E)$ λ-embeds into a tree T, $C = L_j^i \in \mathcal{LP}$ is a cluster of a layering partition of G and v is a vertex of C, then $d_G(v', u) \leq \max\{3\lambda - 1, 2\lambda + 1\}$ for any neighbor $v' \in L^{i-1}$ of v and any $u \in C$.*

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 an arbitrary vertex $u \in C$, $u \neq v$. By the triangle inequality and (1), we have $d_G(a_v, a_u) \leq d_T(a_v, a_u) \leq d_T(a_v, c) + d_T(a_u, c) \leq \lambda$, thus $d_G(a_v, u) \leq d_G(a_v, a_u) + d_G(a_u, u) \leq 2\lambda$. 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 $d_G(v', u) \leq d_G(v', a_v) + d_G(a_v, u) \leq \lambda - 1 + 2\lambda = 3\lambda - 1$, $d_G(v', u) \leq \max\{3\lambda - 1, 2\lambda + 1\}$. $\square$

To make the embedding non-contractive, it suffices to assign the length $\ell := \max\{3\lambda - 1, 2\lambda + 1\}$ to each edge of H and get a uniformly weighted tree $H_\ell = (V, F, \ell)$. Then $d_G(u, v) \leq d_{H_\ell}(u, v) \leq \max\{3\lambda - 1, 2\lambda + 1\}(d_G(u, v) + 2)$. The tree H_ℓ (without Steiner points) provides a 9-approximation to our problem. If we allow Steiner points and assign the length $\ell := \frac{3\lambda}{2}$ to each edge of H', then get a uniformly weighted tree H'_ℓ such that $d_G(u, v) \leq d_{H'_\ell}(u, v) \leq 3\lambda(d_G(u, v) + 1)$.
For a graph $G = (V, E)$, we do not know λ in advance, however we know from Lemma 4 that $\lambda^*(G, \mathcal{T}) \geq D/3$. Therefore, the length ℓ to be assigned to the edges of the tree H (which is defined independently of the value of λ), can be found as follows: $\ell = \max\{d_G(u, v) : uv$ is an edge of $H\}$. The length ℓ, which needs to be assigned to each edge of H', can be found as follows: $\ell = \frac{1}{2}\max\{D, \max\{d_G(u, v) : uv$ is an edge of $H\}\}$. Hence, ℓ can be computed in $O(|V||E|)$ time. Our main result of this section is the following theorem.

Theorem 1. *There exists a factor 6 approximation algorithm for the optimal distortion of embedding an unweighted graph G into a tree.*

The approximation ratio 6 of our algorithm holds only for adjacent vertices of G. It decreases when distances in G increase. 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 [12].

4 Minors, relaxed minors, and metric minors

We define metric relaxed minors, which, together with layering partitions, are used for approximate embedding of graphs into outerplanar metrics.

4.1. Minors and relaxed minors. A graph H is a *minor* of a graph G if a graph isomorphic to H can be obtained from G by contracting or delating some edges and some isolated vertices. To adapt the concept of minor to our embedding purposes, note that $H = (V', E')$ is a minor of $G = (V, E)$ if 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 vertices $v \neq v'$ 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 ends in $G(\mu(u))$ and $G(\mu(v))$;
(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 edges $e = (x, y), e' = (u, v)$ of H, P_e and $P_{e'}$ intersect iff $\{x, y\} \cap \{u, v\} \neq \emptyset$ and if $e = (x, y), e' = (x, v)$, then $P_e \cap P_{e'} = \mu(x)$.

Indeed, if μ exists, then contracting each $\mu(v), v \in V'$, to a single vertex v and each P_e to an edge e, (ii),(iii), and (v') ensure that the resulting graph will be isomorphic to H. Note that if in (v') two paths P_e and $P_{e'}$ intersect, then they intersect in $G(\mu(u))$, where u is the common end of e and e'. In particular, if e, e' are non-incident, then P_e and $P_{e'}$ are disjoint. For our metric purposes we need a weaker notion of minor by allowing intersecting paths to intersect anywhere. A graph $H = (V', E')$ is a *relaxed minor* of a graph $G = (V, E)$ if there exists a map $\mu : V' \cup E' \mapsto 2^V$ satisfying (i)-(iv) and the following relaxation of (v'):

(v) for any two non-incident edges e, e' of H, the paths $P_e \cap P_{e'} = \emptyset$.

The concept of relaxed minor is weaker than that of minor: the triangle C_3 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 path of $K_{1,3}$ between the leaves $\mu(u)$ and $\mu(v)$. The map μ satisfies (i)-(v) but does not satisfy (v'). Relaxed and α-metric relaxed minors (see Subsection 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.* 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.*

4.2. α-Metric relaxed minors. Two sets A, B 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 $G = (V, E)$ if there exists a map $\mu : V' \cup E' \mapsto 2^V$ satisfying (i)-(v) and the following stronger version of condition (v):

(v$^+$) for any 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.

Let φ be an embedding of a graph $G = (V, E)$ into a graph $G' = (V', E')$ with distortion $\leq \alpha$. For $S \subseteq V$ 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 G α-embeds into G' and two sets of vertices A, B inducing connected subgraphs of G are α-far, then $[\varphi(A)] \cap [\varphi(B)] = \emptyset$.*

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

Proof. Suppose 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. Extend φ from V 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 μ and φ map connected subgraphs to connected subgraphs. 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 *for any distinct vertices v, v' of H, $\mu(v)$ and $\mu(v')$ are α-far and for any vertex v and any edge e of H with $v \notin e$, $\mu(v)$ and $\mu(e) = P_e$ are α-far.* We will prove the first part. Since H is 2-connected, any two vertices v, v' belong to a common cycle C of H. Since H is triangle-free, v and v' belong to non-incident edges e, e' of C. Applying (v^+) to e and e', we conclude that $\mu(v)$ and $\mu(v')$ are α-far. Now, we define the following map $\nu : V' \cup E' \mapsto 2^{V(G')}$ from H to G'. For each $v \in V'$, set $\nu(v) = \varphi(\mu(v))$. For each edge $e = uv$ of H, $\mu(e) = P_e$ is a path of 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'. Let $\nu(e)$ be any path of G' between $u' = \varphi(u^*)$ 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 follows that ν satisfies (i) and (iii). We will show that ν satisfies (ii), (iv), and (v) as well. To verify (ii), pick two vertices u, v of H. The sets $\mu(u)$ and $\mu(v)$ are α-far, thus Lemma 6 implies that $\nu(u) = \varphi(\mu(u))$ and $\nu(v) = \varphi(\mu(v))$ are disjoint, showing (ii). Analogously, if v is a vertex and e is an edge of H with $v \notin e$, then, since the sets $\mu(v)$ and $P_e = \mu(e)$ are α-far, thus, by Lemma 6, $\nu(v) = \varphi(\mu(v))$ and $\varphi(P_e)$ are disjoint. Since $\nu(e) \subseteq \varphi(P_e)$, $\nu(v)$ 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, ν satisfies (i)-(v), i.e., H is a relaxed minor of G'. By Proposition 2, H is a minor of G', contradicting that G' is H-minor free. $\square$

4.3. Lower bounds for α-embeddings into $K_{2,r}$-minor free graphs. We use the previous results to give lower bounds for the distortion of embedding a graph $G = (V, E)$ into $K_{2,r}$-minor free graphs.

Proposition 4. *If a cluster C of a layering partition $\mathcal{L}P$ of G contains $r \geq 3$ vertices $v_1^*, \ldots, v_r^*$ 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. Let $\mathcal{LP}$ be defined with respect to s and let T be a BFS tree rooted at s. Let k be the distance from s to C. Since C contains $(4\alpha+2)$-far vertices $v_1^*, \ldots, v_r^*$, $k \geq 2\alpha+2$. We will define a mapping μ from $K_{2,r}$ to G allowing to conclude that $K_{2,r}$ is an α-metric relaxed minor of G. Since $K_{2,r}$ is a subdivided graph, Proposition 3 will show that any embedding of G into a $K_{2,r}$-minor free graph has distortion $> \alpha$.

Let $u_1, \ldots, u_r, v, w$ be the vertices of $K_{2,r}$, where v, w have degree r. Denote by e_i the edge vu_i and by f_i the edge wu_i, $i = 1, \ldots, r$. Let $P_1, \ldots, P_r$ be the paths of T of length $\alpha+1$ from $v_1^*, \ldots, v_r^*$ towards the root s. Denote by $u_1^*, \ldots, u_r^*$ the other end vertices of the paths $P_1, \ldots, P_r$. Let $R_1, \ldots, R_r$ be the paths of T of length $\alpha+1$ from $u_1^*, \ldots, u_r^*$ towards s. Denote by $w_1^*, \ldots, w_r^*$ the other end vertices of the paths $R_1, \ldots, R_r$. Set $\mu(u_i) := u_i^*$, $\mu(e_i) := P_i$ and $\mu(f_i) := R_i$ for $i = 1, \ldots, r$. Let $\mu(v)$ be the connected subgraph of G induced by all (or some) paths connecting the vertices $v_1^*, \ldots, v_r^*$ outside the ball $B_{k-1}(s)$. Finally, let $\mu(w) := B_{k-2\alpha-2}(s)$ (clearly, $w_1^*, \ldots, w_r^*$ belong to $\mu(w)$). From the definitions of μ and $\mathcal{LP}$, we conclude that μ satisfies (i) and (iii). Since $\mu(v) \subseteq \cup_{j \geq k} L^j, \mu(w) = B_{k-2\alpha-2}(s)$, and the vertices $u_1^* = \mu(u_1), \ldots, u_r^* = \mu(u_r)$ belong to $L^{k-\alpha-1}$, the μ-images of the vertices of $K_{2,r}$ are pairwise α-far in G. Analogously, 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 $x \in P_j \cup R_j$ 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 the μ-images of any vertex and any non-incident edge of $K_{2,r}$ are α-far. It remains to show 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 $K_{2,r}$ is an α-metric relaxed minor of G. $\square$

5 Embedding into outerplanar graphs

We present now the algorithm for approximate embedding of graph metrics into into outerplanar metrics.

5.1. The algorithm. Let $G = (V, E)$ be the input graph and let $\mathcal{LP}$ be a layering partition of G. We assume that $\lambda \geq 1$ is so that each cluster C of $\mathcal{LP}$ contains at most two $(4\lambda+2)$-far vertices (otherwise, by Proposition 4, the optimal distortion is larger than λ). Set $\Lambda := 4\lambda + 2$. We call a cluster C *bifocal* if it has two Λ-far vertices c_1 and c_2. In addition, for such cluster C 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 (we will suppose below that c_1 and c_2 form a diametral pair of C). If $\operatorname{diam}(C) \leq \Lambda$ (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 $\operatorname{diam}(C) > 16\lambda + 12$, otherwise,

if $\Lambda < \mathrm{diam}(C) \leq 16\lambda + 12$, then C is a *medium* cluster. An *almost big cluster* is a medium cluster C such that $\mathrm{diam}(C) > 16\lambda + 10$. 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 $\Delta := 8\lambda + 6$. A bifocal cluster C' is *spread* if both cells C_1, C_2 of its father C are adjacent to C'. Given a cluster C at distance k from s and its son C', 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 $\mathcal{F}(C, C')$. We now ready to describe the algorithm.

Algorithm APPROXIMATION BY OUTERPLANAR METRIC
Input: A graph $G = (V, E)$, a layering partition $\mathcal{LP}$ of G, and λ
Output: An outerplanar graph $G' = (V, E')$ or an answer "not"
1. **For** each cluster C of the layering partition $\mathcal{LP}$ **do**
2.   **If** C has two big sons or C is big and has two spread sons, **then return** "not".
3.   **Else for** each son C' of C **do**
4.   Case 1: **If** C' is small, **then** pick 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 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 all vertices of C'.
7.   Case 4: **If** $C' = C'_1 \cup C'_2$, 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.

5.2. Small, medium, and big clusters. We present here without proof several simple properties of clusters of $\mathcal{LP}$.

Lemma 7. *If C is bifocal, then the diameter of each of its cells is $\leq 2\Lambda$.*

Lemma 8. *If C is bifocal and $\mathrm{diam}(C) = d_G(c_1, c_2) > 12\lambda + 6$, then (i) C is $(\mathrm{diam}(C) - 2\Lambda - 1)$-separated, in particular $C_1 \cap C_2 = \emptyset$ and (ii) $\mathrm{diam}(C_1) \leq \Lambda$ and $\mathrm{diam}(C_2) \leq \Lambda$.*
If C is big, then C is $(8\lambda + 8)$-separated and if C is almost big, then C is $(8\lambda + 6)$-separated, whence big and almost big clusters are Λ-separated.
If C is big or almost big, then $\mathrm{diam}(C_1) \leq \Lambda$ and $\mathrm{diam}(C_2) \leq \Lambda$.

Lemma 9. *If C is big, then C has a bifocal spread son C' such that contracting the four cells of C and C' (but preserving the inter-cell edges), we will obtain a $2K_2$.*

Lemma 10. *If C' 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. Big and almost big clusters are spread.*

Lemma 11. *If C is big, no son of C has a cell adjacent to both cells of C. No big cluster C has a small son adjacent to both cells of C.*

5.3. Correctness of the algorithm. The following results establish the correctness and the approximation ratio of our algorithm.

Theorem 2. *Let $G = (V, E)$ be a graph and $\lambda \geq 1$. If the algorithm returns the answer "not", then any embedding of G into a $K_{2,3}$-minor free graph has distortion $> \lambda$. If the algorithm returns the outerplanar graph $G' = (V, E')$, then assigning to its edges weight $w := 20\lambda + 15$, we obtain an embedding of G to G' such that $d_G(x, y) \leq d_{G'}(x, y) \leq 5wd_G(x, y)$ $\forall x, y \in V$. As a result, we obtain a factor $100\lambda + 75$ approximation of the optimal distortion of embedding a graph 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 has distortion $> \lambda$. These conditions hold in the following two cases: (i) 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'. By construction, G' is outerplanar. Let $d_{G'}(x, y)$ be the distance in G' between x and y, where each edge of G' has length $w := 20\lambda + 15$. We continue with the basic property of G' allowing to analyze the approximation ratio.

Proposition 6. *For each edge xy of G, 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 G', $d_G(x, y) \leq 20\lambda + 15$.*

5.4. Proof of Proposition 6. We start with first assertion. First suppose that $d_G(s, x) = d_G(s, y)$. Let C be the cluster of G containing xy. 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 $x \in C, y \in C'$ and C' is a son of C. Let C_0 be 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 C_0, 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, 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 $G : (y, z = c_1, u, c_2, x_{C_0}, x)$.

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, y has a neighbor $x' \in C$. Let $x' \neq x$, otherwise we are done. If C is not big, then $d_G(x, x') \leq 16\lambda + 12$, whence $d_G(x, y) \leq 16\lambda + 13$. So, suppose C is big. If x, 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, let $x \in C_1$ and $x' \in C_2$. By Lemma 11, C' is medium or big. 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 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''. 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. Let a', b' be the vertices of P' at distance $\lambda + 1$ (measured in P') from w', where a' is between w' and x' and b' is between w' and y'. Let L' be the subpath of P' between a' and w' and R' the subpath of P' between w' and b'. Analogously, for P'' we can define the vertices a'', b'' and the paths L'', R'' of length $\lambda + 1$ each. Finally, denote by P_1', P_2' the subpaths of P' between a' and x' and between b' and y'. Analogously, define the supbaths 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.

Now we 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 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 remaining conditions of a λ-metric relaxed minor. The proof of this is subdivided into several results: (1) $d_G(w', C) \geq 4\lambda + 3$ and $d_G(w'', C) \geq 4\lambda + 3$, (2) S is λ-far from all elements of μ except J', J'' (3) w' is λ-far from all elements of μ except L', R' and w'' is λ-far from all elements of μ except L'', R'', (4) L', R' are λ-far from $L'', R'', P_1'', P_2'', J', J''$ and L'', R'' are λ-far from P_1', P_2', J', J'', (5) Q' is λ-far from the R', R'', J'' and Q'' is λ-far L', L'', J', and (6) Q' and Q'' are λ-far.

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 first condition of Proposition 5 is fulfilled. This establishes (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 Λ. 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 (ii).

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 $d_G \leq d_{G'} \leq (100\lambda + 75)d_G$.

References

1. R. Agarwala, V. Bafna, M. Farach, B. Narayanan, M. Paterson, M. Thorup, On the approximability of numerical taxonomy (fitting distances by tree metrics), *SIAM J. Comput.* **28** (1999), 1073–1085.
2. M. Bădoiu, J. Chuzhoy, P. Indyk, A. Sidiropoulos, Low-distortion embeddings of general metrics into the line, *STOC* 2005.
3. M. Bădoiu, E.D. Demaine, M.T. Hajiaghayi, A. Sidiropoulos, M. Zadimoghaddam, Ordinal embedding: approximation algorithms and dimensionality reduction, *APPROX-RANDOM* 2008.
4. M. Bădoiu, P. Indyk, A. Sidiropoulos, Approximation algorithms for embedding general metrics into trees, *SODA* 2007.
5. A. Brandstädt, V. Chepoi, F. Dragan, Distance approximating trees for chordal and dually chordal graphs. *J. Algorithms* **30** (1999).
6. V. Chepoi, F. Dragan, A note on distance approximating trees in graphs, *Europ. J. Combin.* **21** (2000).
7. V. Chepoi, F. Dragan, B. Estellon, M. Habib, Y. Vaxès, Diameters, centers, and approximating trees of δ-hyperbolic geodesic spaces and graphs, *SoCG* 2008.
8. V. Chepoi, B. Fichet, l_∞-Approximation via subdominants, *J. Math. Psychol.* **44** (2000).
9. Y. Dourisboure, F.F. Dragan, C. Gavoille, C. Yan, Spanners for bounded tree-length graphs, *Theor. Comput. Sci.* **383** (2007).
10. Y. Dourisboure, C. Gavoille, Tree-decompositions with bags of small diameter, *Discr. Math.* **307** (2007).
11. Y. Emek, D. Peleg, Approximating minimum max-stretch spanning trees on unweighted graphs, *SODA* 2004.
12. A. Gupta, Steiner points in tree metrics don't (really) help, *SODA* 2001.
13. P. Indyk, J. Matousek, Low-distortion embeddings of finite metric spaces, *Handbook of Discrete and Computational Geometry,* CRC Press LLC, 2004.
14. R. Krauthgamer, J.R. Lee, Algorithms on negatively curved spaces, *FOCS* 2006.
15. Y. Rabinovich, R. Raz, Lower bounds on the distortion of embedding finite metric spaces in graphs *Discr. Comput. Geom.*, **19** (1998).
16. Y. Shavitt, T. Tankel, On internet embedding in hyperbolic spaces for overlay construction and distance estimation, *INFOCOM* 2004.
17. C. Semple, M. Steel, *Phylogenetics*, Oxford University Press, 2003.