

An Improved Bound for Orthogeodesic Point Set Embeddings of Trees

Imre Bárány*

Kevin Buchin†

Michael Hoffmann‡

Anita Liebenau§

Abstract

In an orthogeodesic embedding of a graph, each edge is embedded as an axis-parallel polyline that forms a shortest path in the ℓ_1 metric. In this paper we consider orthogeodesic plane embeddings of trees on grids. A grid is implicitly defined by a set $P \subset \mathbb{R}^2$ of points. Denote by Γ_P the arrangement induced by all horizontal and vertical lines that pass through a point from P . When embedding a graph on the grid defined by P , vertices are mapped to points from P and edges are realized as polylines that bend at vertices of Γ_P only. For integers n and Δ , denote by $t_\Delta(n)$ the minimum number such that for every set P of $t_\Delta(n)$ points in general position, every tree on n vertices with vertex degree at most Δ admits an orthogeodesic plane embedding on the grid defined by P . We show $t_4(n) < 11n/8$ and $t_3(n) < 9n/8$, improving an earlier bound of $3n/2$.

1 Introduction

Given a tree T on n vertices, we want to embed T on an $N \times N$ grid, for some $N \geq n$. In fact, we consider a more restricted setting where possible locations for vertices are specified in form of a set $P \subset \mathbb{R}^2$ of N points. Denote by Γ_P the arrangement induced by all horizontal and vertical lines that pass through a point from P . To embed a graph on the grid defined by P , vertices are mapped to points from P and edges are mapped to arcs that are polylines that bend at vertices of Γ_P only. A point set P is in *general position* if no two points have the same x - or y -coordinate.

A common theme in the study of metric graph embeddings is the desire to control the length of edges. For instance, can every edge be realized as a shortest path? In the Euclidean plane, we arrive at straight line embeddings. A natural counterpart of these embeddings on the grid is called an *orthogeodesic embedding*. In an orthogeodesic embedding, every edge is realized as an *orthogeodesic arc*, that is, a polyline that consists of axis-parallel line segments and forms a shortest path in the ℓ_1 metric. An *L-shaped* arc is an orthogeodesic arc with exactly one bend.

*Renyi Institute of Mathematics and University College London, barany@renyi.hu

†TU Eindhoven, k.a.buchin@tue.nl

‡ETH Zürich, hoffmann@inf.ethz.ch

§Monash University, Anita.Liebenau@monash.edu, work done while at FU Berlin

An embedding is *plane* if no two arcs share a common point that is not a common endpoint. Clearly an orthogeodesic plane embedding can exist only for trees of degree at most four. As it is straightforward to find orthogeodesic embeddings for paths, the only interesting cases are maximum degree three and maximum degree four. For integers n and Δ , denote by $t_\Delta(n)$ the minimum number such that for every set P of $t_\Delta(n)$ points in general position, every tree on n vertices with degree at most Δ admits an orthogeodesic plane embedding on the grid defined by P .

Di Giacomo et al. [2] showed that $t_4(n) \leq 4n-3$ and $t_3(n) \leq 3n/2$. The conference version of [2] (which is reference [1] here) claims that n points are enough for trees of degree at most three. But the proof turned out to be incomplete, as commented in the journal version. Recently, Scheucher [3] showed that $t_4(n) \leq \lfloor (3n-2)/2 \rfloor$. We improve these bounds as follows.

Theorem 1 *For every set $P \subset \mathbb{R}^2$ of $\lfloor (11n-7)/8 \rfloor$ points in general position and every tree T on $n \geq 3$ vertices of degree at most four, T admits an orthogeodesic plane embedding on the grid defined by P .*

Theorem 2 *For every set $P \subset \mathbb{R}^2$ of $\lfloor (9n-4)/8 \rfloor$ points in general position and every tree T on $n \geq 4$ vertices of degree at most three, T admits an orthogeodesic plane embedding on the grid defined by P .*

2 Proofs

For a tree T and $i \in \{1, 2, 3, 4\}$, denote by $d_i(T)$ the number of degree i vertices in T . As a first step we prove Theorem 1 with a weaker bound of $\lfloor (3n-2)/2 \rfloor$ points and a tree T on $n \geq 2$ vertices.

Proof. The idea is to spend one extra point per vertex of degree three or four. Then we need $f(T) := |T| + d_3(T) + d_4(T)$ points. For $n \geq 2$ we have $d_1(T) = 2d_4(T) + d_3(T) + 2$. It follows that $n = |T| = \sum_{i=1}^4 d_i = 3d_4(T) + 2d_3(T) + d_2(T) + 2$ and, therefore, $f(T) = n + (n - 2 - d_2(T) - d_4(T))/2 \leq \lfloor (3n-2)/2 \rfloor$.

We inductively prove the following statement from which the claim follows immediately: For any tree T on $n \geq 1$ vertices, any leaf ℓ of T , any direction d of the four axis directions $\{+x, -x, +y, -y\}$, and any set Γ of $f(T)$ points in general position, there is an orthogeodesic plane embedding of T on the grid defined by Γ such that ℓ is mapped to the extreme point of Γ in direction d and every edge (no edge for

$n = 1$ and one edge, otherwise) connected to ℓ leaves it in the opposite direction. (For instance, if ℓ is mapped to the point λ with largest x -coordinate, then the only incident arc leaves λ on the left side.)

The statement is obviously true for $n \in \{1, 2\}$. For $n \geq 3$ we proceed as follows. Without loss of generality let $d = y$. Map ℓ to the topmost point λ of Γ (there is no choice, anyway), and consider the unique child/neighbor p of ℓ in T . Next subdivide the remaining points of Γ (other than λ) into $\deg_T(p)$ groups. We distinguish three cases depending on $\deg_T(p)$.

Case 1: If $\deg_T(p) \leq 2$, then p is a leaf of $T' = T \setminus \ell$ and we can directly apply induction to T' , y , and $\Gamma \setminus \{\lambda\}$. The edge $\{\ell, p\}$ can be routed going down from λ and then turning left or right to the point π that p is mapped to (Figure 1a).

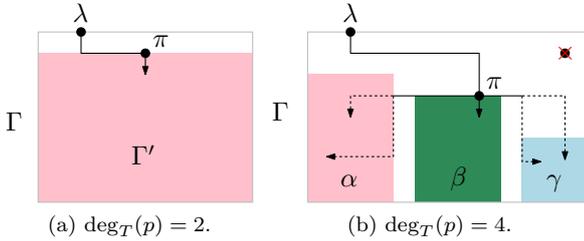


Figure 1: Embedding a vertex p of degree 2 or 4.

Case 2: If $\deg_T(p) = 4$, then consider the tree $T' = T \setminus \ell$. Removal of p splits T' into three components A , B and C . Obtain Γ' from Γ by removing the topmost two points. We partition Γ' into three groups: the leftmost $f(A)$ points go into a set α , the rightmost $f(C)$ points go into a set γ , and the remaining $f(T) - 2 - f(A) - f(C) = |B| + d_3(B) + d_4(B) + 1 = f(B) + 1$ points in between go into a set β . The plan is to embed A on α , $B \cup \{p\}$ on β , and C on γ .

Let π denote the topmost point of β and map p to π . We use the row of Γ below λ to route the edge between λ and π , entering λ from below (as required) and π from above. Now apply induction to three sub-problems, where the vertex p and its corresponding point π are included in all of them. Note that p is a leaf in all of $A' = T' \setminus (B \cup C)$, $B' = T' \setminus (A \cup C)$, and $C' = T' \setminus (A \cup B)$, and that π is the rightmost point of $\alpha' = \alpha \cup \{\pi\}$, the topmost point of $\beta' = \beta$, and the leftmost point of $\gamma' = \gamma \cup \{\pi\}$. (Columns located between the rightmost point of α and π , which belong to β , are ignored for the purpose of handling A' . If π lies above α , also the rows between the topmost row of α and the one of π are ignored. Similarly, some columns and rows are ignored for handling C' . Essentially we always consider square grids.)

Obtain inductively an embedding for A' with p placed in direction x on α' , for B' with p placed in direction y on β' , and for C' with p placed in direction $-x$ on γ' (Figure 1b). The overlay of these three embeddings together with the placement of ℓ at λ and

p at π forms the desired embedding for T : The only edge connected to π in α' enters π from the left side, the only edge connected to π in β' enters π from below, and the only edge connected to π in γ' enters π from the right side. As all edges within each of α' , β' , and γ' are orthogeodesic, no two edges from different sub-problems interfere with each other. (As the dotted lines in Figure 1b suggest, we do not know exactly how π is connected to α and γ . But we *do* know that π is accessed from one particular direction only, and so we can rely on the part shown solid, which is enough to guarantee that these edges do not interfere with those of the embedding of B' .)

Case 3: If $\deg_T(p) = 3$, then without loss of generality let the point of Γ in the row directly below λ be located to the left of λ . Consider the tree $T' = T \setminus \ell$. Removing p from T' splits the tree into two components A and B . Let $A' = T' \setminus B$ and $B' = T' \setminus A$ and partition $\Gamma' = \Gamma \setminus \{\lambda\}$ into two groups: the leftmost $f(A)$ points go into a set α and the remaining $f(T) - 2 - f(A) = |B| + d_3(B) + d_4(B) + 1 = f(B) + 1$ points go into a set β . Denote the topmost point of α by ϕ and embed p at the topmost point π of β .

We distinguish two cases. If π is above ϕ , then by assumption π lies to the left of λ . In this case we route the edge $\{\ell, p\}$ to go down from λ and enter π from the right side (Figure 2a).

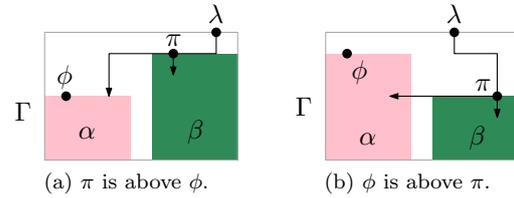


Figure 2: Embedding a degree three vertex at π .

Otherwise, ϕ is above π and we route the edge $\{\ell, p\}$ to go down from λ and turn in the row of ϕ to enter π from above. This is fine, if λ lies to the right of α (Figure 2b and 3b). But if λ lies to the left of β , this routing uses part of the row of ϕ within α , which then is not available for the embedding of A' on $\alpha' = \alpha \cup \{\pi\}$. To be on the safe side, we discard ϕ from Γ (Figure 3a).

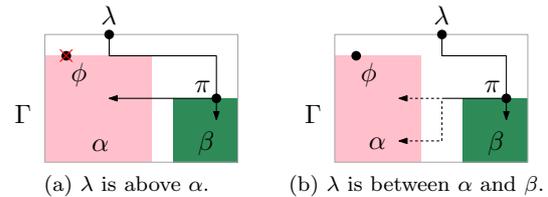


Figure 3: If λ is above α , then ϕ is discarded.

Analogously to Case 2 we inductively obtain embeddings for A' on $\alpha \cup \{\pi\}$ and for B' on β . \square

In order to improve the bound, let us consider the case of a degree three vertex p in the construction more carefully. In Case 3 above we considered three sub-cases and only in the last one (depicted in Figure 3a) a grid point needs to be discarded. If in that case the edge $\{\ell, p\}$ can be routed to leave λ on the left or right side instead of on the bottom side, then ϕ can be kept: If the right side of λ is free, then we route the edge using one bend only (Figure 4a); else if the left side of λ is free, then we move p from B to A and move the leftmost point in β to α in order to map p to ϕ instead and also route the edge $\{\ell, p\}$ using one bend only (Figure 4b).

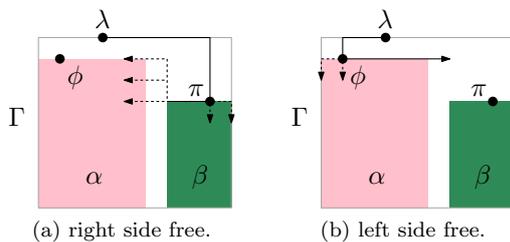


Figure 4: Avoid discarding ϕ if one side of λ is free.

If ℓ has degree at most two in the original tree, then obviously at least one side of λ is free. If ℓ has degree four in the original tree, then obviously no side of λ is free. The situation is less clear in case that ℓ has degree three in the original tree.

In the case depicted in Figure 2a, no side of π is free for embedding on β . But as far as α is concerned, the top side of π may be regarded as free. Although no orthogeodesic path from π to any point in α can leave π on its top side, conceptually the top side is free regardless. This point of view makes sense, because in that case λ is located at a corner of Γ and we may also regard λ as an extreme point on the (say) right side. The assignment of α and β can then be done correspondingly with respect to the right side of Γ (Figure 5) and π can be nicely accessed from the right.

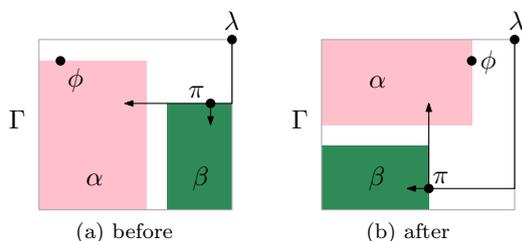


Figure 5: Switching sides if λ is at a corner.

Similarly, it can be checked that in all cases other than the one depicted in Figure 2a both α and β can access λ from at least (typically exactly) one other side. For instance, in the case depicted in Figure 3a the embedding on β can access π from the right,

whereas the embedding on α can access π from above: a geodesic path cannot actually enter π from above, but it can leave α on its right side in any row above π and then move down to the row of π once it reaches the area above β to finally enter π on its left side (Figure 6a); from the perspective of α —which ignores all columns in between its right side and π —that looks like entering π from above. The path can be routed as described, unless π is the leftmost point and, hence, top-left corner of β (Figure 6b). But then we can treat π as leftmost point of β and consider the top side of π free as far as β is concerned, whereas the bottom and left side are both accessible for α .

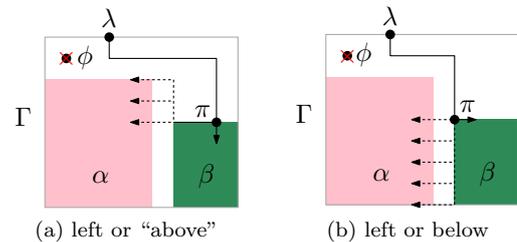


Figure 6: Possible ways for α to access π .

So during our recursive construction we discard a grid point for each degree four vertex and for certain degree three vertices. Consider a fixed processing order defined by a starting leaf of the tree. We partition the degree three vertices into two classes: For a *good* vertex v we guarantee that the left or the right side of the starting point λ is free when processing v . By the analysis above no point needs to be discarded when processing a good degree three vertex. In contrast, for a *bad* vertex no such guarantee holds. For every bad vertex v in a subtree we make an extra point available that can be discarded when processing v . Our goal is to derive a lower bound for the number of vertices that we may regard as good.

Proposition 3 *The parent of a bad degree 3 vertex is a degree 4 vertex or a good degree 3 vertex.*

Proof. Consider a bad degree 3 vertex and let p be its parent. Clearly a child of a vertex of degree at most two is good. Therefore it remains to exclude that p is a bad degree 3 vertex. When processing a bad degree 3 vertex p , we discard the point ϕ and ensure that all children of p are good (cf. Figure 6). \square

Proposition 4 *If a child of a degree 3 vertex p is a bad degree 3 vertex, then the other child of v is a good degree 3 vertex.*

Proof. There is only one case where we cannot guarantee a free side at the starting point for a child of a degree three vertex: in Figure 2a, for the child of p in B to be embedded on β . If the other child of p does

not have degree three, then we select A to contain the degree three child c of p , making c good. \square

Proposition 5 *If a degree 3 vertex v has a child c so that the subtree rooted at c is a path, then v is good.*

Proof. When handling v , we let A be the subtree rooted at c , to be embedded on the left point set α . The only problematic case is the one depicted in Figure 3a, where we have to show how to avoid discarding ϕ . Given that A is a path, it can be embedded on α in a monotone fashion, from right to left: every arc leaves the parent on the left and enters the child along the y -direction. In particular, the part of the row of ϕ to the right of ϕ (that is used by the arc between λ and π) is not touched. \square

Propositions 3–5 allow us to classify degree three vertices during a top-down traversal of T as follows. Initially, all vertices which Proposition 5 applies to are *good*, and the remaining vertices are *unclassified*. When encountering a degree four vertex, all its unclassified degree three children are *bad*. When encountering a good degree three vertex for which both children are unclassified degree three vertices, one of the children is *bad* and the other is *good*. In all other cases, an unclassified degree three child is *good*.

Proof. (of Theorem 1) Let $T = (V, E)$ and denote $d_i = d_i(T)$. Observe that $d_1 = 2d_4 + d_3 + 2$. Let V_3^- denote the set of bad degree three vertices in T , let V_3^+ denote the set of good degree three vertices in T , and let $d_3^- = |V_3^-|$ and $d_3^+ = |V_3^+|$. Let $W \subseteq V$ denote the set of all vertices $v \in V$ such that either $v \in V_3^-$ or v is a leaf. Denote by $F \subseteq E$ the set of all edges in T that are incident to at least one vertex from W . By Propositions 3 and 5 and given $n \geq 3$, every edge in F is incident to exactly one vertex from W and one vertex from $V \setminus W$. Therefore, we can double count by the endpoints in W to obtain $|F| = 3d_3^- + d_1 = 2d_4 + 4d_3^- + d_3^+ + 2$ and by the endpoints in $V \setminus W$ to obtain $|F| \leq 4d_4 + 3d_3^+ + 2d_2$. Combining both bounds we get $2d_3^- \leq d_4 + d_3^+ + d_2 - 1$.

Setting $k = d_4 + d_3^-$, we use $n + k$ grid points for $n = 3d_4 + 2d_3 + d_2 + 2$ vertices. Therefore

$$\begin{aligned} k &= \frac{3}{8} \left(n - \frac{1}{3}d_4 + \frac{2}{3}d_3^- - 2d_3^+ - d_2 - 2 \right) \\ &\leq \frac{3}{8} \left(n - \frac{5}{3}d_3^+ - \frac{2}{3}d_2 - \frac{7}{3} \right) \leq \frac{3n-7}{8}. \quad \square \end{aligned}$$

Proof. (of Theorem 2) Define d_i , V_3^- , V_3^+ , d_3^- and d_3^+ as above. If $d_3 = 0$, then T is a path and the statement is trivial. Hence suppose $d_3 \geq 1$, which implies $n \geq 4$. We consider T as a directed tree by orienting all edges away from the root. Let F denote the set of directed edges (u, v) in T such that $u \in V_3^-$ and $v \in V_3^+$. We claim that $d_3^- + |F| \leq d_3^+ - 1$.

To prove this claim, define an injective map $g: V_3^- \cup F \rightarrow V_3^+$. For a vertex $u \in V_3^-$ let $g(u)$ be the sibling of u in T . Such a sibling $g(u)$ exists by Proposition 3 and $g(u) \in V_3^+$ by Proposition 4. As a vertex in V_3^+ has at most one sibling, g is injective on V_3^- .

For an edge $e = (u, v) \in F$ set $g(e) = v$, where $v \in V_3^+$ by definition. As every vertex in V_3^+ has exactly one incoming edge, g is injective on F .

Note that for all vertices in $g(V_3^-)$ the parent is in V_3^+ , whereas for all vertices in $g(F)$ the parent is in V_3^- . Therefore, g is injective on $V_3^- \cup F$, as claimed. It also follows that the highest vertex in V_3^+ (closest to the root) is not in $g(V_3^- \cup F)$ and, therefore, $|g(V_3^- \cup F)| \leq d_3^+ - 1$.

The claim directly generalizes to the case where F is the set of directed paths (v_0, \dots, v_k) in T such that $k \geq 1$, $v_0 \in V_3^-$, $v_k \in V_3^+$, and $\deg_T(v_i) = 2$, for $0 < i < k$. Then Propositions 3 and 5 imply $|F| = 2d_3^-$. In combination with the claim we obtain $3d_3^- \leq d_3^+ - 1$. We use $n + d_3^-$ grid points for $n = 2d_3 + d_2 + 2$ vertices. It follows that

$$d_3^- = \frac{1}{8} (n + 6d_3^- - 2d_3^+ - d_2 - 2) \leq \frac{1}{8} (n - 4). \quad \square$$

3 Conclusions

As an obvious open problem it remains to tighten the bounds for $t_3(n)$ and $t_4(n)$. Most notably, it would be nice to prove or disprove $t_3(n) = n$. No non-trivial lower bound seems to be known, even for maximum degree four and if the embedding is restricted to use L-shaped arcs only.

Acknowledgments This work was started at the 10th Gremo Workshop on Open Problems in Bergün, Switzerland, June 4–8, 2012. We thank Yoshio Okamoto and Bettina Speckmann for independently bringing the problem investigated to our attention. We also thank all participants of the workshop for the inspiring and productive atmosphere. Finally, we thank one of the reviewers for helpful comments.

References

- [1] E. Di Giacomo, F. Frati, R. Fulek, L. Grilli, and M. Krug. Orthogeodesic point-set embedding of trees. *Graph Drawing (Proc. GD '11)*, volume 7034 of *Lecture Notes Comput. Sci.*, pages 52–63, 2012.
- [2] E. Di Giacomo, F. Frati, R. Fulek, L. Grilli, and M. Krug. Orthogeodesic point-set embedding of trees. *Comput. Geom. Theory Appl.*, 46(8):929–944, 2013.
- [3] M. Scheucher. Orthogeodesic point set embeddings of outerplanar graphs. Master's thesis, Inst. Softw. Tech., TU Graz, Graz, Austria, 2015.