An Algorithm to Construct Greedy Drawings of Triangulations*

Patrizio Angelini[1], Fabrizio Frati[1], and Luca Grilli[2]

[1] Dipartimento di Informatica e Automazione - Roma Tre University
{angelini,frati}@dia.uniroma3.it
[2] Dipartimento di Ingegneria Elettronica e dell'Informazione - Perugia University
luca.grilli@diei.unipg.it

Abstract. We show an algorithm to construct greedy drawings of every given triangulation.

1 Introduction

In a *greedy routing* setting, a node forwards packets to a neighbor that is *closer* to the destination's geographic location. Different distance metrics define different meanings for the word "closer", and consequently define different routing algorithms for the packet delivery. The most used and studied metric is of course the *Euclidean distance*.

The efficiency of the greedy routing algorithms strongly relies on the geographic coordinates of the nodes. This is a drawback of such algorithms, for the following reasons: (i) Nodes of the network have to know their locations, hence they have to be equipped with GPS devices, which are expensive and increase the energy consumption of the nodes; (ii) geographic coordinates are independent of the network obstructions, i.e. obstacles making the communication between two close nodes impossible, and, more in general, they are independent of the network topology; this could lead to situations in which the communication fails because a *void* has been reached, i.e., the packet has reached a node whose neighbors are all farther from the destination than the node itself.

A brilliant solution to such weaknesses has been proposed by Rao *et al.* who in [9] proposed a scheme in which nodes decide *virtual coordinates* and then apply the greedy routing algorithm relying on such coordinates rather than on the real geographic ones. Since virtual coordinates do not need to reflect the nodes actual positions, they can be suitably chosen to guarantee that the greedy routing algorithm delivers packets with high probability. Experiments have shown that such an approach strongly improves the reliability of greedy routing [9,8]. Further, it has been proved that virtual coordinates guarantee greedy routing to work for every connected topology when they can be chosen in the hyperbolic plane [5], and that some modifications of the routing algorithm guarantee that Euclidean virtual coordinates can be chosen so that the packet delivery always succeeds [1], even if the coordinates need to be locally computed [2].

* Work partially supported by MUR under Project "MAINSTREAM: Algoritmi per strutture informative di grandi dimensioni e data streams" and by the Italian Ministry of Research, Grant number RBIP06BZW8, project FIRB "Advanced tracking system in intermodal freight transportation".

I.G. Tollis and M. Patrignani (Eds.): GD 2008, LNCS 5417, pp. 26–37, 2009.

Subsequent to the Rao *et al.* paper [9], an intense research effort has been devoted to determine on which network topologies the Euclidean greedy routing with virtual coordinates is guaranteed to work. From a graph-theoretic point of view, the problem is as follows: Which are the graphs that admit a *greedy embedding*, i.e., a straight-line drawing Γ such that, for every pair of nodes u and v, there exists a *distance-decreasing path* in Γ? A path $(v_0, v_1, \ldots, v_m)$ is distance-decreasing if $d(v_i, v_m) < d(v_{i-1}, v_m)$, for $i = 1, \ldots, m$. In [8] Papadimitriou and Ratajczak conjectured the following:

Conjecture 1. (Papadimitriou and Ratajczak [8]) Every triconnected planar graph admits a greedy embedding.

Papadimitriou and Ratajczak showed that $K_{k,5k+1}$ has no greedy embedding, for $k \geq 1$. As a consequence, both the triconnectivity and the planarity are necessary, because there exist planar non-triconnected graphs, such as $K_{2,11}$, and non-planar triconnected graphs, such as $K_{3,16}$, that do not admit any greedy embedding. Further, they observed that, if a graph G has a greedy embedding, then any graph containing G as a spanning subgraph has a greedy embedding. It follows that Conjecture 1 extends to all graphs which are spanned by a triconnected planar graph. Related to such an observation, they proved that every triconnected graph not containing a $K_{3,3}$-minor has a triconnected planar spanning subgraph.

For a few classes of triconnected planar graphs the conjecture is easily shown to be true, for example graphs with a *Hamiltonian path* and *Delaunay Triangulations*. At SODA'08 [3], Dhandapani proved the conjecture for the first non-trivial class of triconnected planar graphs, namely he showed that every *triangulation* admits a greedy embedding. The proof of Dhandapani is probabilistic, namely the author proves that among all the *Schnyder drawings* of a triangulation [10], there exists a drawing which is greedy. Although such a proof is elegant, relying at the same time on an old Combinatorial Geometry theorem, known as the *Knaster-Kuratowski-Mazurkievicz Theorem* [6], and on standard Graph Drawing techniques, as the *Schnyder realizers* [10] and the *canonical orderings* of a triangulation [4], it does not lead to an embedding algorithm.

In this paper we show an algorithm for constructing greedy drawings of triangulations. The algorithm relies on a different and maybe more intuitive approach with respect to the one used in [3]. We define a simple class of graphs, called *binary cactuses*, and we provide an algorithm to construct a greedy drawing of any binary cactus. Finally, we show how to find, for every triangulation, a binary cactus spanning it. It is clear that the previous statements imply an algorithm for constructing greedy drawings of triangulations. Namely, consider any triangulation G, apply the algorithm to find a binary cactus S spanning G, and then apply the algorithm to construct a greedy drawing of S. As already observed, adding edges to a greedy drawing leaves the drawing greedy, hence S can be augmented to G, obtaining the desired greedy drawing of G.

Theorem 1. *Given a triangulation G, there exists an algorithm to compute a greedy drawing of G.*

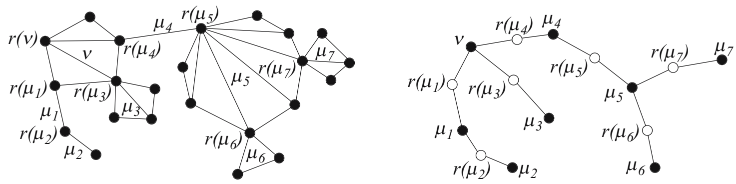

(a) (b)

Fig. 1. (a) A binary cactus S. (b) The block-cutvertex tree of S. White (resp. black) circles represent C-nodes (resp. B-nodes).

2 Preliminaries

A graph G is *connected* if every pair of vertices of G is connected by a path. A *cutvertex* is a vertex whose removal increases the number of connected components of G. A connected graph is *biconnected* if it has no cutvertices. The maximal biconnected subgraphs of a graph are its *blocks*. Each edge of G falls into a single block of G, while cutvertices are shared by different blocks. The *block-cutvertex tree*, or BC-tree, of a connected graph G is a tree with a B-node for each block of G and a C-node for each cutvertex of G. Edges in the BC-tree connect each B-node μ to the C-nodes associated with the cutvertices in the block of μ.

The BC-tree of G may be thought as rooted at a specific block ν. When the BC-tree $\mathcal{T}$ of a graph G is rooted at a certain block ν, we denote by $G(\mu)$ the subgraph of G induced by all vertices in the blocks contained in the subtree of $\mathcal{T}$ rooted at μ. In a rooted BC-tree $\mathcal{T}$ of a graph G, for each B-node μ we denote by $r(\mu)$ the cutvertex of G parent of μ in $\mathcal{T}$. If μ is the root of $\mathcal{T}$, i.e., $\mu = \nu$, then we let $r(\mu)$ denote any non-cutvertex node of the block associated with μ. In the following, unless otherwise specified, each considered BC-tree is meant to be rooted at a certain B-node ν such that the block associated with ν has at least one vertex $r(\nu)$ which is not a cutvertex. It is not difficult to see that such a block exists in every planar graph.

A *rooted triangulated binary cactus* S, in the following simply called *binary cactus*, is a connected graph such that (see Fig 1): (i) the block associated with each B-node of $\mathcal{T}$ is either an edge or a *triangulated cycle*, i.e., a cycle $(r(\mu), u_1, u_2, \ldots, u_h)$ triangulated by the edges from $r(\mu)$ to each of $u_1, u_2, \ldots, u_h$; (ii) every cutvertex is shared by exactly two blocks of S.

A *planar drawing* of a graph is a mapping of each vertex to a distinct point of the plane and of each edge to a Jordan curve between its endpoints such that no two edges intersect except, possibly, at common endpoints. A planar drawing of a graph determines a circular ordering of the edges incident to each vertex. Two drawings of the same graph are *equivalent* if they determine the same circular ordering around each vertex. A *planar embedding* is an equivalence class of planar drawings. A planar drawing partitions the plane into topologically connected regions, called *faces*. The unbounded face is the *outer face*. The outer face of a graph G is denoted by $f(G)$. A *chord* of a graph G is an edge connecting two non-adjacent vertices of $f(G)$. A graph together with a planar embedding and a choice for its outer face is called *plane graph*. A plane graph is a *triangulation* when all its faces are triangles. A plane graph is *internally-triangulated*

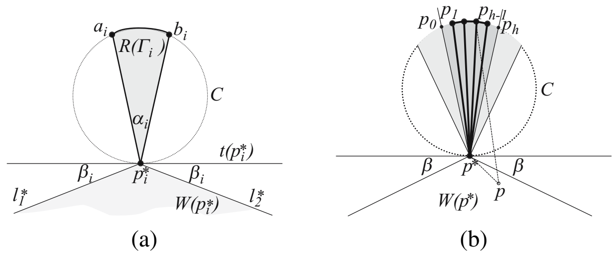

Fig. 2. (a) Illustration for Properties 1–3 of Γ. (b) Base case of the algorithm. The light and dark shaded region represents $R(\Gamma)$ (the angle of $R(\Gamma)$ at p^* is α). The dark shaded region represents the intersection of $W(p^*, \alpha/2)$ with the circle delimited by C.

when all its internal faces are triangles. An *outerplane* graph is a plane graph such that all its vertices are incident to the outer face. A *Hamiltonian cycle* of a graph G is a simple cycle passing through all vertices of G. Notice that a biconnected outerplane graph has only one Hamiltonian cycle, the one delimiting its outer face.

3 Greedy Drawing of a Binary Cactus

In this section, we give an algorithm to compute a greedy drawing of a binary cactus S. Such a drawing is constructed by a bottom-up traversal of the BC-tree $\mathcal{T}$ of S.

Consider the root μ of a subtree of $\mathcal{T}$ corresponding to a block of S, consider the k children of μ, which correspond to cutvertices of S, and consider the children of such cutvertices, say $\mu_1, \mu_2, \dots, \mu_k$. Notice that each C-node child of μ is parent of exactly one B-node μ_i of $\mathcal{T}$, by definition of binary cactus. For each $i = 1, \dots, k$, inductively assume to have a drawing Γ_i of $S(\mu_i)$ satisfying the following properties. Let α_i and β_i be any two angles less than $\pi/4$ such that $\beta_i \geq \alpha_i$. Refer to Fig. 2.a.

- *Property 1.* Γ_i is a greedy drawing.
- *Property 2.* Γ_i is entirely contained inside a region $R(\Gamma_i)$ delimited by an arc (a_i, b_i) of a circumference C and by two segments (p_i^*, a_i) and (p_i^*, b_i), such that p_i^* is a point of C and the diameter through p_i^* cuts (a_i, b_i) in two arcs of the same length. The angle $\widehat{a_i p_i^* b_i}$ is α_i.
- *Property 3.* Consider the tangent $t(p_i^*)$ to C in p_i^*. Consider two half-lines l_1^* and l_2^* incident to p_i^*, lying on the opposite part of C with respect to $t(p_i^*)$, and forming angles equal to β_i with $t(p_i^*)$. Denote by $W(p_i^*)$ the wedge centered at p_i^*, delimited by l_1^* and l_2^*, and not containing C. Then, for every vertex v in $S(\mu_i)$ and for every point p internal to $W(p_i^*)$, a distance-decreasing path $(v = v_0, v_1, \dots, v_l = r(\mu_i))$ from v to $r(\mu_i)$ exists in Γ_i such that $d(v_j, p) < d(v_{j-1}, p)$ for $j = 1, \dots, l$.

In the base case, block μ has no child. Denote by $(r(\mu) = u_0, u_1, \dots, u_{h-1})$ the block of S corresponding to μ. If $h = 2$, i.e., μ corresponds to an edge, draw such an edge as a vertical segment, with u_1 above u_0. A region $R(\Gamma_i)$ can be easily constructed, for every angles α and β, with $\beta \geq \alpha$, satisfying the above properties. If $h > 2$, i.e., μ corresponds to a triangulated cycle of S, place $r(\mu)$ at any point p^* and consider a wedge $W(p^*, \alpha/2)$ that has an angle equal to $\alpha/2$, that is incident to $r(\mu)$, and that is bisected by the vertical half-line incident to $r(\mu)$ and directed upward (see Fig. 2.b). Denote by p'_a and p'_b the intersection points of the half-lines delimiting $W(p^*, \alpha/2)$ with a circumference C through $r(\mu)$, properly intersecting the border of $W(p^*, \alpha/2)$ twice. Denote by A the arc of C between p'_a and p'_b not containing p^*. Consider points $p'_a = p_0, p_1, \ldots, p_h = p'_b$ on A such that the distance between any two consecutive points p_i and p_{i+1} is the same. Place vertex u_i at point p_i, for $i = 1, 2, \ldots, h-1$.

We show that the constructed drawing Γ satisfies Property 1. Consider any two vertices u_i and u_j, with $i < j$. If $i = 0$, then u_0 and u_j are joined by an edge, which provides a distance-decreasing path among them. Otherwise, we claim that path $(u_i, u_{i+1}, \ldots, u_j)$ is distance-decreasing. In fact, for each $l = i, i+1, \ldots, j-2$, angle $\widehat{u_l u_{l+1} u_j}$ is greater than $\pi/2$, because triangle (u_l, u_{l+1}, u_j) is inscribed in less than half a circumference with u_{l+1} as middle point. Hence, (u_l, u_j) is the longest side of triangle (u_l, u_{l+1}, u_j) and $d(u_{l+1}, u_j) < d(u_l, u_j)$ follows. Drawing Γ satisfies Property 2 by construction. In order to prove that Γ satisfies Property 3, we have to show that, for every vertex u_i, with $i \geq 1$, and for every point p in $W(p^*)$, $d(u_0, p) < d(u_i, p)$. However, angle $\widehat{pp^*p_i}$ is at least $\beta + (\frac{\pi}{2} - \frac{\alpha}{4})$, which is more than $\pi/2$. It follows that segment $\overline{pp_i}$ is the longest side of triangle (p, p^*, p_i), thus proving that $d(u_0, p) < d(u_i, p)$.

Now suppose μ is a node of $\mathcal{T}$ having k children. We show how to construct a drawing Γ of $S(\mu)$ satisfying Properties 1–3 with parameters α and β. Denote by $(r(\mu) = u_0, u_1, \ldots, u_{h-1})$ the block of S corresponding to μ. Consider any circumference C with center c. Let p^* be the point of C with smallest y-coordinate. Consider wedges $W(p^*, \alpha)$ and $W(p^*, \alpha/2)$ with angles α and $\alpha/2$, respectively, incident to p^* and such that the diameter of C through p^* is their bisector. Region $R(\Gamma)$ is the intersection region of $W(p^*, \alpha)$ with the closed circle delimited by C.

Consider a circumference C' with center c intersecting the two lines delimiting $W(p^*, \alpha/2)$ in two points p'_a and p'_b such that angle $\widehat{p'_a c p'_b} = 3\alpha/2$. Denote by p' the intersection point between C' and (c, p^*). Observe that angle $\widehat{p'_a p' p'_b} = 3\alpha/4$. Denote by A the arc of C' delimited by p'_a and p'_b not containing p'. Consider points $p'_a = p_0, p_1, \ldots, p_h = p'_b$ on A such that the distance between any two consecutive points p_i and p_{i+1} is the same. Observe that, for each $i = 0, 1, \ldots, h-1$, angle $\widehat{p_i c p_{i+1}} = \frac{3\alpha}{2h}$.

First, we draw the block of S corresponding to μ. As in the base case, place vertex $u_0 = r(\mu)$ at p^* and, for $i = 1, 2, \ldots, h-1$, place u_i at point p_i. Recursively construct a drawing Γ_i of $S(\mu_i)$ satisfying Properties 1–3 with $\alpha_i = \frac{3\alpha}{16h}$ and $\beta_i = \frac{3\alpha}{8h}$.

We are going to place each drawing Γ_i of $S(\mu_i)$ together with the drawing of the block of S corresponding to μ, thus obtaining a drawing Γ of $S(\mu)$. Not all h nodes u_i are cutvertices of S. However, with a slight abuse of notation, we suppose that block $S(\mu_i)$ has to be placed at node u_i. Refer to Fig 3. Consider point p_i and its "neighbors" p_{i-1} and p_{i+1}. Consider lines $t(p_{i-1})$ and $t(p_{i+1})$ tangent to C' through p_{i-1} and p_{i+1}, respectively. Further, consider circumferences C_{i-1} and C_{i+1} centered at p_{i-1} and p_{i+1}, respectively, and passing through p_i. Moreover, consider lines h_{i-1} and h_{i+1} through p_i and tangent to C_{i-1} and C_{i+1}, respectively. For each point p_i, consider two half-lines t_1^i and t_2^i incident to p_i, cutting C' twice, and forming angles $\beta_i = \frac{3\alpha}{8h}$ with $t(p_i)$. Denote by $W(p_i)$ the wedge delimited by t_1^i and t_2^i and containing c.

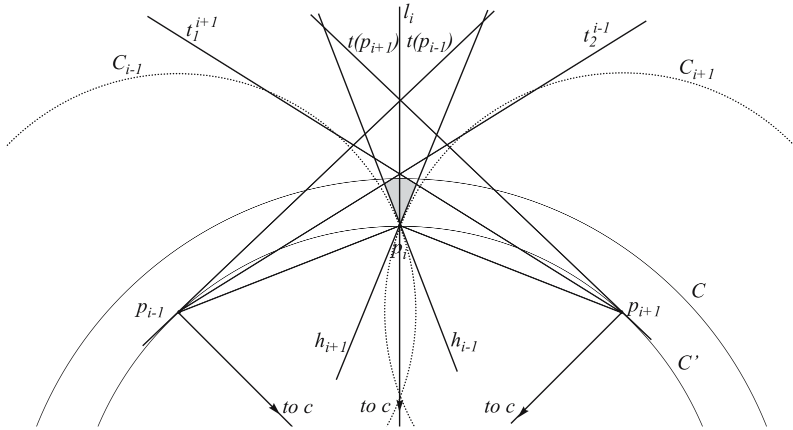

Fig. 3. Lines and circumferences in the construction of Γ. The shaded region is R_i.

We place Γ_i inside the bounded region R_i intersection of the half-plane H^{i-1} delimited by h_{i-1} and not containing C_{i-1}, of the half-plane H^{i+1} delimited by h_{i+1} and not containing C_{i+1}, of $W(p_{i-1})$, of $W(p_{i+1})$, and of the circle delimited by C.

First, we show that R_i is "large enough" to contain Γ_i, namely we claim that there exists an isosceles triangle T that has an angle larger than $\alpha_i = \frac{3\alpha}{16h}$ incident to p_i and that is completely contained in R_i. Such a triangle will have the further feature that the angle incident to p_i is bisected by the half-line l_i incident to c and passing through p_i.

Lines h_{i-1} and h_{i+1} are both passing through p_i; we prove that they have different slopes and we compute the angles they form at p_i. Line h_{i-1} forms an angle of $\pi/2$ with segment $\overline{p_{i-1}p_i}$; angle $\widehat{cp_ip_{i-1}}$ is equal to $\frac{\pi}{2} - \frac{3\alpha}{4h}$, since $\widehat{p_icp_{i-1}} = \frac{3\alpha}{2h}$ and since triangle (p_{i-1}, c, p_i) is isosceles. Hence, the angle delimited by h_{i-1} and l_i is $\pi - \pi/2 - (\frac{\pi}{2} - \frac{3\alpha}{4h}) = \frac{3\alpha}{4h}$. Analogously, the angle between l_i and h_{i+1} is $\frac{3\alpha}{4h}$. Hence, the intersection of H^{i-1} and H^{i+1} is a wedge $W(p_i, h_{i-1}, h_{i+1})$ centered at p_i, with an angle of $\frac{3\alpha}{2h}$, and bisected by l_i. We claim that each of t_2^{i-1} and t_1^{i+1} cuts the border of $W(p_i, h_{i-1}, h_{i+1})$ twice. The angle between $t(p_{i-1})$ and $\overline{p_{i-1}p_i}$ is $\frac{3\alpha}{4h}$, namely the angle between $t(p_{i-1})$ and $\overline{cp_{i-1}}$ is $\pi/2$, and angle $\widehat{cp_{i-1}p_i}$ is $\frac{\pi}{2} - \frac{3\alpha}{4h}$. The angle between $t(p_{i-1})$ and t_2^{i-1} is $\beta_i = \frac{3\alpha}{8h}$ by construction. Hence, the angle between t_2^{i-1} and $\overline{p_{i-1}p_i}$ is $\frac{3\alpha}{4h} - \frac{3\alpha}{8h} = \frac{3\alpha}{8h}$. Since the slope of both h_{i-1} and h_{i+1} with respect to $\overline{p_{i-1}p_i}$ is greater than $\frac{3\alpha}{8h}$ and less than $\pi - \frac{3\alpha}{8h}$, namely the slope of h_{i-1} and h_{i+1} with respect to $\overline{p_{i-1}p_i}$ is $\frac{\pi}{2}$ and $\frac{\pi}{2} + \frac{3\alpha}{2h}$, respectively (notice that $\alpha \leq \pi/4$ and $h \geq 2$), then t_2^{i-1} intersects both h_{i-1} and h_{i+1}. It can be analogously proved that t_1^{i+1} intersects h_{i-1} and h_{i+1}. It follows that the intersection of H^{i-1}, H^{i+1}, $W(p_{i-1})$, and $W(p_{i+1})$ contains a triangle T as required by the claim (the angle of T incident to p_i is $\frac{3\alpha}{2h}$). Considering circumference C does not invalidate the existence of T, since C is concentric with C' and has a bigger radius, hence T can be chosen sufficiently small so that it completely lies inside C.

Now Γ_i can be placed inside T, by scaling Γ_i down till it fits inside T (see Fig. 4.a). The scaling always allows Γ_i to be placed inside T, since the angle of $R(\Gamma_i)$ incident to p is $\alpha_i = \frac{3\alpha}{16h}$, that is smaller than the angle of T incident to p_i, which is $\frac{3\alpha}{2h}$. In

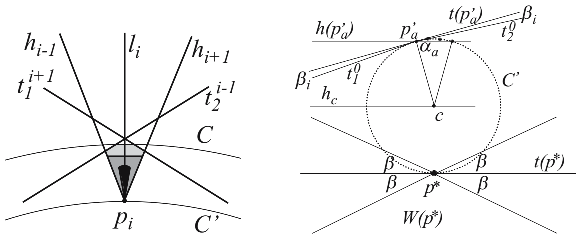

(a) (b)

Fig. 4. (a) Placement of Γ inside R_i. Region $R(\Gamma)$ is the darkest, triangle T is composed of $R(\Gamma)$ and of the second darkest region, R_i is composed of T and of the light shaded region. (b) Illustration for the proof of Lemma 1.

particular, we choose to place Γ_i inside T so that l_i bisects the angle of $R(\Gamma_i)$ incident to p_i. This concludes the construction of Γ. We have the following lemmata.

Lemma 1. *The closed wedge $W(p^*)$ is completely contained inside the open wedge $W(p_i)$, for each $i = 0, 1, \ldots, h$.*

Proof: Consider any point p_i. Observe that p_i is contained inside the wedge $\overline{W}(p^*)$ obtained by reflecting $W(p^*)$ with respect to $t(p^*)$. Namely, p_i is contained inside $W(p^*, \alpha/2)$, which is in turn contained inside $\overline{W}(p^*)$, since $\alpha/2 < \pi - 2\beta$, as a consequence of the fact that $\pi/4 > \beta \geq \alpha$. Hence, in order to prove the lemma, it suffices to show that the absolute value of the slope of each of t_1^i and t_2^i is less than the absolute value of the slope of the half-lines delimiting $W(p^*)$. Such latter half-lines form angles of β, by construction, with the x-axis.

The slope of t_1^i can be computed by summing up the slope of t_1^i with respect to $t(p_i)$ with the slope of $t(p_i)$. The former slope is equal to $\beta_i = \frac{3\alpha}{8h}$, by construction. Recalling that $t(p_i)$ is the tangent to A in p_i, the slope of $t(p_i)$ is bounded by the maximum among the slopes of the tangents to points of A. Such a maximum is clearly achieved at p'_a and p'_b and is equal to $3\alpha/4$. Namely, refer to Fig. 4.b and consider the horizontal lines $h(c)$ and $h(p'_a)$ through c and p'_a, respectively, that are traversed by radius (c, p'_a). Such a radius forms angles of $\pi/2$ with $t(p'_a)$; hence, the slope of $t(p'_a)$, that is equal to the angle between $t(p'_a)$ and $h(p'_a)$, is $\pi/2$ minus the angle α_a between $h(p'_a)$ and (c, p'_a). Angle α_a is the alternate interior of the angle between $h(c)$ and (c, p'_a), which is complementary to the half of angle $\widehat{p'_a c p'_b}$, which is equal to $3\alpha/2$, by construction. It follows that α_a is equal to $\frac{\pi}{2} - \frac{3\alpha}{4}$ and the slope of $t(p'_a)$ is $\frac{3\alpha}{4}$.

Hence, the slope of t_1^i is at most $\frac{3\alpha}{4} + \frac{3\alpha}{8h}$, which is less than α, since $h \geq 2$, and hence less than β. Analogously, the slope of t_2^i is less than β, and the lemma follows. □

Corollary 1. *Point p^* is inside the open wedge $W(p_i)$, for each $i = 1, 2, \ldots, h$.*

Lemma 2. *For every pair of indices i and j such that $1 \leq i < j \leq k$, the drawing of $S(\mu_j)$ is contained inside $W(p_i)$ and the drawing of $S(\mu_i)$ is contained inside $W(p_j)$.*

Proof: If $S(\mu_i)$ and $S(\mu_j)$ are consecutive, i.e., the cutvertices parents of $S(\mu_i)$ and $S(\mu_j)$ are u_i and u_j and $j = i + 1$, then the statement is true by construction. Suppose

$S(\mu_i)$ and $S(\mu_j)$ are not consecutive. Consider the triangle T_i delimited by (p^*, p_i), by t_2^i, and by the line through p^* and p'_b. T_i contains the triangle delimited by (p^*, p_{i+1}), by t_2^{i+1}, and by the line through p^* and p'_b, which in turn contains the triangle delimited by (p^*, p_{i+2}), by t_2^{i+2}, and by the line through p^* and p'_b. Repeating such an argument shows that T_i contains the triangle T_{j-1} delimited by (p^*, p_{j-1}), by t_2^{j-1}, and by the line through p^* and p'_b. By construction, Γ_j lies inside T_{j-1}, and the lemma follows. □

We prove that the constructed drawing Γ satisfies Properties 1–3.

Property 1. We show that, for every pair of vertices w_1 and w_2, there exists a distance-decreasing path between them in Γ. If both w_1 and w_2 are internal to the same graph $S(\mu_i)$, the property follows by induction. If one of w_1 and w_2, say w_1, is $r(\mu)$ and the other one, say w_2, is a node in $S(\mu_i)$ then, by Property 3, there exists a distance-decreasing path $(w_2 = v_0, v_1, \ldots, v_l = r(\mu_i))$ from w_2 to $r(\mu_i)$ such that, for every point p in $W(p_i)$, $d(v_j, p) < d(v_{j-1}, p)$, for $j = 1, 2, \ldots, l$. By Corollary 1, p^* is contained inside $W(p_i)$. Hence path $(w_2 = v_0, v_1, \ldots, v_l = r(\mu_i), w_1 = r(\mu))$ is a distance-decreasing path between w_1 and w_2. If w_1 belongs to $S(\mu_i)$ (possibly $w_1 = u_i$) and w_2 belongs to $S(\mu_j)$ (possibly $w_2 = u_j$) then suppose, w.l.o.g., that $j > i$. We show the existence of a distance-decreasing path $\mathcal{P}$ in Γ, composed of three subpaths $\mathcal{P}_1, \mathcal{P}_2$, and $\mathcal{P}_3$. By Property 3, Γ_j is such that there exists a distance-decreasing path $\mathcal{P}_1 = (w_1 = v_0, v_1, \ldots, v_l = r(\mu_i))$ from w_1 to $r(\mu_i)$ such that, for every point p in $W(p_i)$, $d(v_j, p) < d(v_{j-1}, p)$, for $j = 1, 2, \ldots, l$. By Lemma 2, drawing Γ_j, and hence vertex w_2, is contained inside $W(p_i)$, hence path $\mathcal{P}_1$ decreases the distance from w_2 at every vertex. Path $\mathcal{P}_2 = (u_i = r(\mu_i), u_{i+1}, \ldots, u_j = r(\mu_j))$ is easily shown to decrease the distance from w_2 at every vertex. In fact, for each $l = i, i+1, \ldots, j-2$, angle $\widehat{u_l u_{l+1} u_j}$ is greater than $\pi/2$, because triangle (u_l, u_{l+1}, u_j) is inscribed in less than half a circumference with u_{l+1} as middle point. Angle $\widehat{u_l u_{l+1} w_2}$ is strictly greater than $\widehat{u_l u_{l+1} u_j}$, hence it is the biggest angle in triangle (u_l, u_{l+1}, w_2) and $d(u_{l+1}, w_2) < d(u_l, w_2)$ follows. By induction, there exists a distance-decreasing path $\mathcal{P}_3$ from $r(\mu_j)$ to w_2, thus obtaining a distance-decreasing path $\mathcal{P}$ from w_1 to w_2.

Property 2. Such a property holds for Γ by construction.

Property 3. Consider any node v in $S(\mu_i)$ and any point p internal to $W(p^*)$. By Lemma 1, p is internal to $W(p_i)$. By induction, there exists a distance-decreasing path $(v = v_0, v_1, \ldots, v_l = r(\mu_i))$ such that $d(v_j, p) < d(v_{j-1}, p)$, for $j = 1, 2, \ldots, l$. Hence, path $(v = v_0, v_1, \ldots, v_l = r(\mu_i), v_{l+1} = r(\mu))$ is a distance-decreasing path such that $d(v_j, p) < d(v_{j-1}, p)$, for $j = 1, 2, \ldots, l+1$, if and only if $d(r(\mu), p) < d(r(\mu_i), p)$. However, angle $\widehat{pr(\mu)r(\mu_i)}$ is at least $\beta + (\frac{pi}{2} - \frac{3\alpha}{8})$, which is more than $\pi/2$. Hence, $(p, r(\mu_i))$ is the longest side of triangle $(p, r(\mu), r(\mu_i))$, thus proving that $d(r(\mu), p) < d(r(\mu_i), p)$, and Property 3 holds for Γ.

When the induction on $\mathcal{T}$ is performed with $\mu = \nu$, we obtain a greedy drawing of S, thus proving the following:

Theorem 2. *There exists an algorithm that constructs a greedy drawing of any binary cactus.*

4 Spanning a Triangulation with a Binary Cactus

In this section we prove the following theorem:

Theorem 3. *Given a triangulation G, there exists a spanning subgraph S of G such that S is a binary cactus.*

Consider any triangulation G. We are going to construct a binary cactus S spanning G. First, we outline the algorithm to construct S. Such an algorithm has several steps. At the first step, we choose a vertex u incident to $f(G)$ and we construct a triangulated cycle C_T composed of u and all its neighbors. We remove u and its incident edges from G, obtaining a biconnected internally-triangulated plane graph G^*. At the beginning of each step after the first one, we suppose to have already constructed a binary cactus S whose vertices are a subset of the vertices of G (at the beginning of the second step, S coincides with C_T), and to have a set $\mathcal{G}$ of subgraphs of G (at the beginning of the second step, G^* is the only graph in $\mathcal{G}$). Each of such subgraphs is biconnected, internally-triangulated, has an outer face whose vertices already belong to S, and has internal vertices. All such internal vertices do not belong to S and each vertex of G not belonging to S is internal to a graph in $\mathcal{G}$. Only one of the graphs in $\mathcal{G}$ may have chords (at the beginning of the second step, G^* is such a graph). During each step, we perform the following two actions: (1) We partition the only graph G_C of $\mathcal{G}$ with chords, if any, into several biconnected internally-triangulated chordless plane graphs; we remove G_C from $\mathcal{G}$ and we add to $\mathcal{G}$ all graphs with internal vertices into which G_C has been partitioned; (2) we choose a graph G_i from $\mathcal{G}$, we choose a vertex u incident to the outer face of G_i and already belonging to exactly one block of S, and we add to S a block composed of u and of all its neighbors internal to G_i. We remove u and its incident edges from G_i, obtaining a biconnected internally-triangulated plane graph G_i^*. We remove G_i from $\mathcal{G}$ and we add G_i^* to $\mathcal{G}$. The algorithm stops when $\mathcal{G}$ is empty.

Now we give the details of the above outlined algorithm. At the first step of the algorithm, choose any vertex u incident to $f(G)$. Consider the neighbors $(u_1, u_2, \dots, u_l)$ of u in clockwise order around it. Since G is a triangulation, $C = (u, u_1, u_2, \dots, u_l)$ is a cycle. Let C_T be the triangulated cycle obtained by adding to C the edges connecting u to its neighbors. Let $S = C_T$. Remove vertex u and its incident edges from G, obtaining a biconnected internally-triangulated graph G^*. If G^* has no internal vertex, then all the vertices of G belong to S and we have a binary cactus spanning G. Otherwise, let $\mathcal{G} = \{G^*\}$. For each graph $G_i \in \mathcal{G}$, consider the vertices incident to $f(G_i)$. Each of such vertices can be either *forbidden for* G_i or *assigned to* G_i. A vertex w is forbidden for G_i if the choice of not introducing in S any new block incident to w and spanning a subgraph of G_i has been done. Conversely, a vertex w is assigned to G_i if a new block incident to w and spanning a subgraph of G_i could be introduced in S. For example, w is forbidden for G_i if there exist two blocks of S sharing w as a cutvertex. At the end of the first step of the algorithm, choose any two vertices incident to $f(G^*)$ as the only forbidden vertices for G^*. All other vertices incident to $f(G^*)$ are assigned to G^*. At the beginning of the i-th step, with $i \geq 2$, we assume that each of the following holds:

- *Invariant A:* Graph S is a binary cactus spanning all and only the vertices that are not internal to any graph in $\mathcal{G}$.

- *Invariant B:* Each graph in $\mathcal{G}$ is biconnected, internally-triangulated, and has internal vertices.
- *Invariant C:* Only one of the graphs in $\mathcal{G}$ may have chords.
- *Invariant D:* No internal vertex of a graph $G_i \in \mathcal{G}$ belongs to a graph $G_j \in \mathcal{G}$.
- *Invariant E:* For each graph $G_i \in \mathcal{G}$, all vertices incident to $f(G_i)$ are assigned to G_i, except for two vertices, which are forbidden.
- *Invariant F:* Each vertex v incident to the outer face of a graph in $\mathcal{G}$ is assigned to at most one graph $G_v \in \mathcal{G}$. The same vertex is forbidden for all graphs $\overline{G}_v \in \mathcal{G}$ such that v is incident to $f(\overline{G}_v)$ and $\overline{G}_v \neq G_v$.
- *Invariant G:* Each vertex assigned to a graph in $\mathcal{G}$ belongs to exactly one block of S.

Such invariants clearly hold after the first step of the algorithm.

Action 1: If all graphs in $\mathcal{G}$ are chordless, go to Action 2. Otherwise, by Invariant C, only one of the graphs in $\mathcal{G}$, say G_C, may have chords. We use such chords to partition G_C into k biconnected, internally-triangulated, chordless graphs G_C^j, with $j = 1, 2, \dots, k$. Consider the biconnected outerplane subgraph O_C of G_C induced by the vertices incident to $f(G_C)$. To each internal face f of O_C delimited by a cycle c, a graph G_C^j is associated such that G_C^j is the subgraph of G_C induced by the vertices of c or inside c. Before replacing G_C with graphs G_C^j in $\mathcal{G}$, we show how to decide which vertices incident to the outer face of a graph G_C^j are assigned to G_C^j and which vertices are forbidden for G_C^j. Since each graph G_C^j is univocally associated with a face of O_C (namely the face of O_C delimited by the cycle that delimits $f(G_C^j)$), in the following we assign vertices to the faces of O_C and we forbid vertices for the faces of O_C, meaning that if a vertex is assigned to (forbidden for) a face of O_C delimited by a cycle c then it is assigned to (resp. forbidden for) graph G_C^j whose outer face is delimited by c.

We want to assign the vertices incident to $f(O_C)$ to faces of O_C so that the following properties are satisfied. *Property 1:* No forbidden vertex is assigned to any face of O_C. *Property 2:* No vertex is assigned to more than one face of O_C; *Property 3:* Each face of O_C has exactly two incident vertices which are forbidden for it; all other vertices of the face are assigned to it.

By Invariant E, G_C has two forbidden vertices. We construct an assignment of vertices to faces of O_C in some steps. Let p be the number of chords of O_C. Consider the Hamiltonian cycle O_C^0 of O_C, and assign all vertices of O_C^0, but for the two forbidden vertices, to the only internal face of O_C^0. At the i-th step, $1 \leq i \leq p$, we insert into O_C^{i-1} a chord of O_C, obtaining a graph O_C^i. This is done so that Properties 1–3 are satisfied by O_C^i (with O_C^i instead of O_C). After all p chords of O_C have been inserted, $O_C^p = O_C$, and we have an assignment of vertices to faces of O_C satisfying Properties 1–3. Properties 1–3 are clearly satisfied by the assignment of vertices to faces of O_C^0. Inductively assume Properties 1–3 are satisfied by the assignment of vertices to faces of O_C^{i-1}. Let (u_a, u_b) be the chord that is inserted at the i-th step. Chord (u_a, u_b) partitions a face f of O_C^{i-1} into two faces f_1 and f_2. By Property 3, two vertices u_1^* and u_2^* incident to f are forbidden for it and all other vertices incident to f are assigned to it. For each face of O_C^i different from f_1 and f_2, assign and forbid vertices as in the same face in O_C^{i-1}. Assign and forbid vertices for f_1 and f_2 as follows.

– If vertices u_a and u_b are the same vertices of u_1^* and u_2^*, assign to each of f_1 and f_2 all vertices incident to it, except for u_a and u_b. No forbidden vertex has been assigned to any face of O_C^i (Property 1). Vertices u_a and u_b have not been assigned to any face. All vertices assigned to f belong to exactly one of f_1 and f_2 and so they have been assigned to exactly one face (Property 2). The only vertices of f_1 (resp. of f_2) not assigned to it are u_a and u_b, while all other vertices are assigned to such a face (Property 3).
– If vertices u_a and u_b are both distinct from u_1^* and u_2^* and both u_1^* and u_2^* are in the same of f_1 and f_2, say in f_1, assign to f_1 all vertices incident to it, except for u_1^* and u_2^*, and assign to f_2 all vertices incident to it, except for u_a and u_b. No forbidden vertex has been assigned to any face of O_C^i (Property 1). Vertices u_a and u_b have been assigned to exactly one face. All other vertices assigned to f belong to exactly one of f_1 and f_2 and so they have been assigned to exactly one face (Property 2). The only vertices of f_1 (resp. of f_2) not assigned to it are u_1^* and u_2^* (resp. u_a and u_b), while all other vertices are assigned to such a face (Property 3).
– If vertices u_a and u_b are both distinct from u_1^* and u_2^* and one of u_1^* and u_2^*, say u_1^*, is in f_1 while the other one, say u_2^*, is in f_2, assign to f_1 all vertices incident to it, except for u_1^* and u_a, and assign to f_2 all vertices incident to it, except for u_2^* and u_b. No forbidden vertex has been assigned to any face of O_C^i (Property 1). Vertices u_a and u_b have been assigned to exactly one face. All other vertices assigned to f belong to exactly one of f_1 and f_2 and so they have been assigned to exactly one face (Property 2). The only vertices of f_1 (resp. of f_2) not assigned to it are u_1^* and u_a (resp. u_2^* and u_b), while all other vertices are assigned to such a face (Property 3).
– If one of vertices u_1^* and u_2^* coincides with one of u_a and u_b, say u_1^* coincides with u_a, and the other one, say u_2^*, is in one of f_1 and f_2, say in f_1, assign to f_1 all vertices incident to it, except for u_2^* and u_a, and assign to f_2 all vertices incident to it, except for u_a and u_b. No forbidden vertex has been assigned to any face of O_C^i (Property 1). Vertex u_a has not been assigned to any face and vertex u_b has been assigned to exactly one face. All other vertices assigned to f belong to exactly one of f_1 and f_2 and so they have been assigned to exactly one face (Property 2). The only vertices of f_1 (resp. of f_2) not assigned to it are u_2^* and u_a (resp. u_a and u_b), while all other vertices are assigned to such a face (Property 3).

Graph G_C is removed from $\mathcal{G}$. All graphs G_C^j having internal vertices are added to $\mathcal{G}$. It is easy to see that Invariants A–G are satisfied after Action 1.

Action 2: After Action 1 all graphs in $\mathcal{G}$ are chordless. There is at least one graph G_i in $\mathcal{G}$, otherwise the algorithm would have stopped before Action 1. By Invariant B, G_i has internal vertices. Choose any vertex u incident to $f(G_i)$ and assigned to G_i. Since G_i is biconnected and has internal vertices, $f(G_i)$ has at least three vertices. Since each graph in $\mathcal{G}$ has at most two forbidden vertices (by Invariant E), a vertex u assigned to G_i exists. Consider all the neighbors $(u_1, u_2, \dots, u_l)$ of u internal to G_i, in clockwise order around u. Since G is biconnected, chordless, internally triangulated, and has internal vertices, then $l \geq 1$. If $l = 1$ then let C_T be edge (u, u_1). Otherwise, let C_T be the triangulated cycle obtained by adding to cycle $(u, u_1, u_2, \dots, u_l)$ the edges connecting u to its neighbors. Add C_T to S. Remove u and its incident edges from G_i, obtaining a graph G_i^*. Assign to G_i^* all vertices incident to $f(G_i^*)$, except for the two vertices forbidden for G_i. Remove G_i from $\mathcal{G}$ and insert G_i^*, if it has internal vertices, in $\mathcal{G}$. It is easy to see that Invariants A–G are satisfied after Action 2.

When the algorithm stops, i.e., when there is no graph in $\mathcal{G}$, by Invariant A graph S is a binary cactus spanning all vertices of G, hence proving Theorem 3.

5 Conclusions

In this paper we have shown an algorithm for constructing greedy drawings of triangulations. The algorithm relies on two main results. The first one states that every binary cactus admits a greedy drawing. The second result, that may be of its own interest, is that, for every triangulation G, there exists a binary cactus S spanning G.

After this paper was submitted, the authors realized that a slight modification of the two main arguments, presented in Sect. 3 and 4, proves Conjecture 1. Namely, it can be shown that every triconnected planar graph can be spanned by a rooted *non-triangulated* binary cactus, i.e. a connected graph such that the block associated with each B-node of $\mathcal{T}$ is either an edge or a cycle and every cutvertex is shared by exactly two blocks. A greedy drawing of such a graph can be constructed by the drawing algorithm presented for rooted *triangulated* binary cactuses (the proof that the drawings constructed by the algorithm are greedy is slightly more involved due to the absence of edges $(r(\mu), u_i)$, for $i = 2, 3, \cdots, h-2$). However, two reviewers of our paper made us aware that the conjecture has been positively settled by Leighton and Moitra in a paper to appear at FOCS'08 [7]. The approach used by Leighton and Moitra is surprisingly similar to ours.

References

1. Ben-Chen, M., Gotsman, C., Gortler, S.J.: Routing with guaranteed delivery on virtual coordinates. In: CCCG 2006 (2006)
2. Ben-Chen, M., Gotsman, C., Wormser, C.: Distributed computation of virtual coordinates. In: Erickson, J. (ed.) SoCG 2007, pp. 210–219. ACM Press, New York (2007)
3. Dhandapani, R.: Greedy drawings of triangulations. In: Huang, S.-T. (ed.) SODA 2008, pp. 102–111. SIAM, Philadelphia (2008)
4. de Fraysseix, H., Pach, J., Pollack, R.: How to draw a planar graph on a grid. Combinatorica 10(1), 41–51 (1990)
5. Kleinberg, R.: Geographic routing using hyperbolic space. In: INFOCOM 2007, pp. 1902–1909. IEEE, Los Alamitos (2007)
6. Knaster, B., Kuratowski, C., Mazurkiewicz, C.: Ein beweis des fixpunktsatzes fur n dimensionale simplexe. Fundamenta Mathematicae 14, 132–137 (1929)
7. Leighton, T., Moitra, A.: Some results on greedy embeddings in metric spaces. In: FOCS 2008 (2008)
8. Papadimitriou, C.H., Ratajczak, D.: On a conjecture related to geometric routing. Theor. Comput. Sci. 344(1), 3–14 (2005)
9. Rao, A., Papadimitriou, C.H., Shenker, S., Stoica, I.: Geographic routing without location information. In: Johnson, D.B., Joseph, A.D., Vaidya, N.H. (eds.) MOBICOM 2003, pp. 96–108. ACM Press, New York (2003)
10. Schnyder, W.: Embedding planar graphs on the grid. In: SODA 1990, pp. 138–148. SIAM, Philadelphia (1990)