



On a conjecture related to geometric routing

Christos H. Papadimitriou¹, David Ratajczak*

Computer Science Division, University of California at Berkeley, Berkeley, CA 94720-1776, USA

Abstract

We conjecture that any planar 3-connected graph can be embedded in the plane in such a way that for any nodes s and t , there is a path from s to t such that the Euclidean distance to t decreases monotonically along the path. A consequence of this conjecture would be that in any ad hoc network containing such a graph as a spanning subgraph, two-dimensional virtual coordinates for the nodes can be found for which the method of purely greedy geographic routing is guaranteed to work. We discuss this conjecture and its equivalent forms show that its hypothesis is as weak as possible, and show a result delimiting the applicability of our approach: any 3-connected $K_{3,3}$ -free graph has a planar 3-connected spanning subgraph. We also present two alternative versions of greedy routing on virtual coordinates that provably work. Using Steinitz's theorem we show that any 3-connected planar graph can be embedded in three dimensions so that greedy routing works, albeit with a modified notion of distance; we present experimental evidence that this scheme can be implemented effectively in practice. We also present a simple but provably robust version of greedy routing that works for any graph with a 3-connected planar spanning subgraph.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Geometric routing; Ad hoc networks; Planar graphs; Greedy routing; Convex embeddings; Planar embeddings

1. Introduction

Routing in the internet is typically based on internet protocol (IP) addresses, a hierarchical address space that enables swift and effective routing decisions. Recently, research

* Corresponding author.

E-mail addresses: christos@cs.berkeley.edu (C.H. Papadimitriou), dratajcz@cs.berkeley.edu (D. Ratajczak).

¹ Supported by NSF ITR grant CCR-0121555.

in networking has increasingly focused on ad hoc networks [2,13], consisting of sensors communicating via wireless antennae and acting both as origins, destinations, and routers of messages. However, there remains no universally accepted system of addresses or routing in such networks. *Geometric routing* [1,3,6] is a family of routing algorithms using the geographic coordinates of the nodes as addresses for the purpose of routing. One such algorithm is *Euclidean greedy routing*, which is attractive for its simplicity; each node forwards the packet to the neighbor that has the closest Euclidean distance to the destination address. Unfortunately, purely greedy routing sometimes fails to deliver a packet because of the phenomenon of *voids* or “lakes” (nodes with no neighbor closer to the destination). To recover from failures of greedy routing, various forms of *face routing* have been proposed [4,3], in which the presence of a void triggers a special routing mode until the greedy mode can be reestablished. Several schemes have been proposed [1,6] that are guaranteed to deliver a message with certain performance guarantees.

Geometric routing is complicated by two factors. First, since GPS antennae are relatively costly both in price and energy consumption, it is unlikely that ad hoc networks in the foreseeable future can rely on the availability of precise geographic coordinates. Second, the precise coordinates may be disadvantageous as they do not account for obstructions or other topological properties of the network. To address these concerns, Rao et al. [12] recently proposed a scheme in which the nodes first decide on fictitious *virtual coordinates*, and then apply greedy routing based on those. The coordinates are found by a distributed version of the *rubber band* algorithm originally due to Tutte [15] and used often in graph theory [8]. It was noted, on the basis of extensive experimentation, that this approach makes greedy routing much more reliable (in Section 3.1 of this paper we present experimental results on a slight variant of that scheme that has even better performance). However, despite the solid grounding of the ideas in [12] in geometric graph theory, no theoretical results are known for such schemes, and a rigorous examination of virtual coordinates has only recently been initiated [5,10]. The present paper is an attempt to fill this gap: we use sophisticated ideas from geometric graph theory in order to prove the existence of sound virtual coordinate routing schemes.

The focal point of this paper is a novel conjecture in geometric graph theory that is elegant, plausible, has important consequences, and seems to be deep. Consider the embedded graph shown in Fig. 1. It has the following property: given any two distinct nodes s and t , there is a neighbor of s that is closer in Euclidean distance to t than s is. We call such an embedding a *greedy embedding*. We conjecture that *any planar, 3-connected graph has a greedy embedding*. Since every such graph has a convex planar embedding [15], one in which all faces are convex, they are natural candidates for our conjecture, even though there are certainly other graphs with greedy embeddings (for example any graph with a Hamiltonian path has a greedy embedding on a straight line). Furthermore, since the existence of a greedy embedding is a monotonic graph property, in that adding an edge cannot deprive a graph of greedy embeddability, the conjecture extends immediately to any graph with a 3-connected planar subgraph spanning the vertices.

While we have been unable to prove the conjecture, we present several interesting, related results: we show that it is tight, in that both planarity and 3-connectivity are necessary. We present a family of examples that suggest that its proof must rely on local arguments. We prove in Theorem 2 that any graphs without a $K_{3,3}$ minor must contain a planar 3-connected

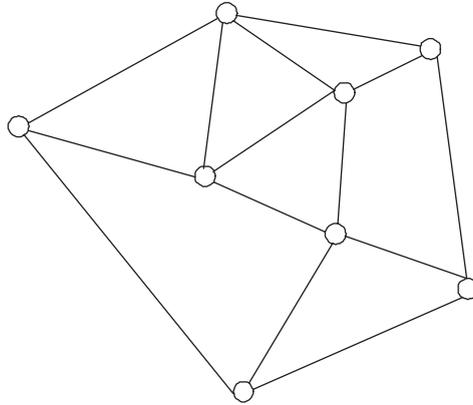


Fig. 1. A greedy embedding.

subgraph and therefore, by our conjecture, greedy routing works for such graphs. We present an alternative greedy routing algorithm that works for all planar 3-connected graphs that is based on a three-dimensional embedding. Finally, we present a simple face routing algorithm that is guaranteed (Theorem 4) to work on the embeddings obtained by the rubber band algorithm [12] on 3-connected planar graphs.

2. The conjecture

Let $G = (V, E)$ be a graph. An *embedding of G in the plane* is a one-to-one map \mathbf{e} from V to \mathbb{R}^2 , where \mathbf{e} maps each edge of the graph to the line segment joining the images of its endpoints; the embedding is *planar* when no two such segments intersect at any point other than their endpoints.

For a given graph G and embedding \mathbf{e} , we write $d(u, v)$, where $u, v \in E$ and \mathbf{e} is implicit, to denote the Euclidean distance of the images of u and v . We say that a path (v_0, v_1, \dots, v_m) is *distance decreasing* if $d(v_i, v_m) < d(v_{i-1}, v_m)$ for $i = 1, \dots, m$.

We propose the following conjecture:

Conjecture 1. *Any planar, 3-connected graph can be embedded on the plane so that between any two vertices s and t there is a distance-decreasing path beginning with s and ending with t .*

By *greedy routing algorithm*, we mean the following non-deterministic algorithm which, given two nodes s (the origin) and t (the destination), and a graph embedded in the plane, produces a path (v_0, v_1, \dots) as follows:

```

set  $i = 0, v_0 = s$ 
while  $v_i$  is not  $t$  do
  if exists  $u$  adjacent to  $v_i$  such that  $d(v_i, t) > d(u, t)$ 
    set  $i = i + 1$  and  $v_i = u$ 
  else stop

```

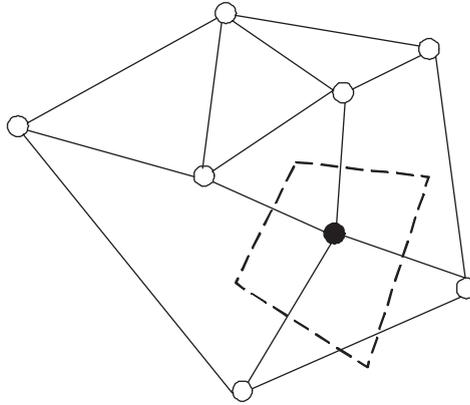


Fig. 2. The cell of a node v is all points in the plane with v as the closest node.

We call an embedding *greedy* if this algorithm always terminates with $v_i = t$. We shall show that the embeddings conjectured above are precisely the greedy embeddings.

To restate the greedy property in a slightly more geometric form, consider an embedding of a graph, a node v of that embedding, and all of its neighbors u_1, \dots, u_k . The *cell* of v is the set of all points in the plane that are closer to v than to any u_i (see Fig. 2). Note that the cell of a node may be unbounded.

Theorem 1. *The following are equivalent statements about a planar embedding of a graph:*

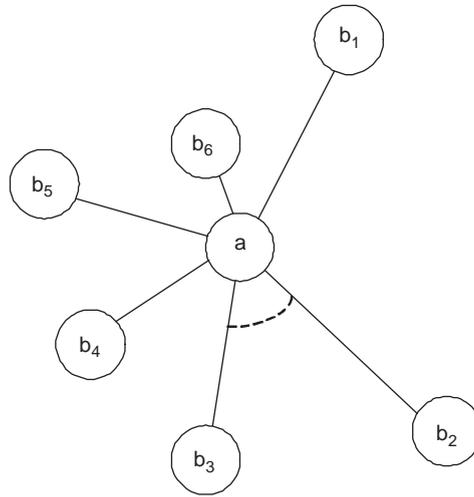
- (1) *The embedding is greedy.*
- (2) *Between any two vertices s and t there is a distance-decreasing path beginning with s and ending with t .*
- (3) *For any two nodes s and t with $s \neq t$ there is a neighbor r of s such that $d(r, t) < d(s, t)$.*
- (4) *The cell of any node v contains no other vertex than v itself.*

Proof.

- (1 \Rightarrow 2) The greedy algorithm returns such a path for any s and t .
- (2 \Rightarrow 3) The first step of the distance-decreasing path yields such a neighbor.
- (3 \Rightarrow 1) The greedy algorithm from s to t must terminate at some node t' , since the distances to t decrease at each step. If $t' \neq t$, this means that t' has no neighbor closer to t , contradicting (3).
- (3 \Rightarrow 4) If the cell of v contains a node u , then u is not a neighbor of v and any neighbor of v has a greater distance to u than v does, contradicting (3).
- (4 \Rightarrow 3) If (3) does not hold, then t is closer to s than it is to any of the neighbors of s , implying t is in the cell of s . \square

2.1. Some examples

The following simple lemma is handy in constructing graphs with no greedy embedding:

Fig. 3. $K_{1,6}$ has no greedy embedding.

Lemma 1. *In a greedy embedding, any node t must have an edge to the closest node u in the embedding.*

Proof. Otherwise, u has no neighbor that is closer to t than itself. \square

We may therefore construct, for any $k > 0$, a k -connected graph that has no greedy embedding.

Proposition 1. $K_{k,5k+1}$ has no greedy embedding for $k > 0$.

Proof. Let A denote the set of nodes in the partition of size k and B denote the remaining nodes. In any embedding, for each element of B we identify which element of A is closest to it. By the pigeonhole principle, there must be 6 nodes $b_1, \dots, b_6 \in B$ that have some $a \in A$ as the closest element of A . Consider the angles formed by $b_i a b_{i+1 \pmod{6}}$, as illustrated in Fig. 3. At least one of these angles must be no greater than $\pi/3$. Suppose this angle is between b_i and b_{i+1} . By the law of sines, one of the edges ab_i or ab_{i+1} must be no shorter than $b_i b_{i+1}$. This means that one of b_i or b_{i+1} has no edge to the node that is closest to it in the embedding, which by Lemma 1 implies that the embedding is not greedy. \square

It follows from the proposition that the hypotheses of the conjecture are necessary, in that there exist counterexamples to the conjecture that are planar but not 3-connected ($K_{2,11}$), or 3-connected but not planar ($K_{3,16}$).

Taking a different track, the following gives an interesting *sufficient* condition for an embedding to be greedy:

Proposition 2. *If a convex embedding has no angle $2\pi/3$ or larger between consecutive edges on a face, then it is greedy.*

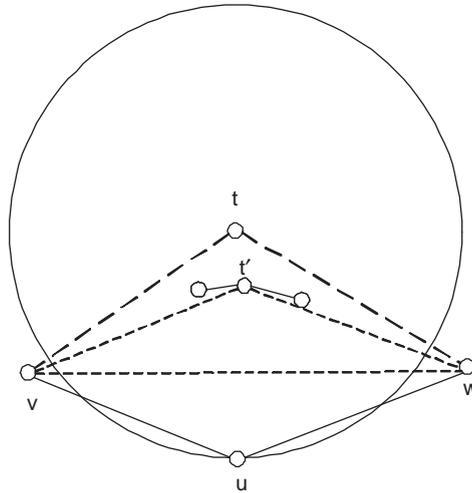


Fig. 4. A non-greedy embedding has a big angle.

Proof. We shall prove the contrapositive. Assume that an embedding is not greedy; we shall exhibit an angle larger than $2\pi/3$. Since the embedding is not greedy there are two nodes u and t such that no neighbor of u is closer to t than u is. Consider the consecutive neighbors of u , v and w , such that t is in the region formed by angle vuw ; by our assumption, v and w do not lie in the interior of the circle with center t and radius $|tu|$ (see Fig. 4).

Consider the node t' within triangle vtw that is closest to the side vw (it could be t itself). The interior of the quadrilateral $uvt'w$ contains no vertex, since no node lies within triangle vuw (by convexity) and such a vertex would lie closer to vw than t' . Thus, there is an angle formed by consecutive edges incident upon t' that is at least as large as angle $vt'w$ —which is at least as large as angle wtv .

We conclude that there is an angle in the embedding that is at least as large as angle wtv . However, because v and w are not in the interior of the circle, it follows that $2 \cdot vuw + vtw \geq 2\pi$, which implies that one of the two angles (vuw or the angle of t') is at least $2\pi/3$. \square

Proposition 2 is of limited applicability in proving the conjecture, since any face of 6 or more sides must contain an angle at least $2\pi/3$. In fact, even triangulated graphs—such as the *flower graph*, obtained by tiling the plane with the pattern shown in Fig. 5—can require an angle greater than $2\pi/3$.

The flower graph does have a greedy embedding. Consider the nodes whose obtuse angles are represented with bold lines in Fig. 5. For greedy routing to work between these nodes, the diamond region in bold must be skewed slightly. We can do this by shrinking the small equilateral triangles down to negligible points, and then translating these points slightly in the manner illustrated in Fig. 6. From this example, we see that the proof of the conjecture cannot simply rely on the existence of an appropriate global embedding, in the style of Tutte's proof [15] or Andre'ev's [14], but would require breaking symmetries in the graph based on local criteria.

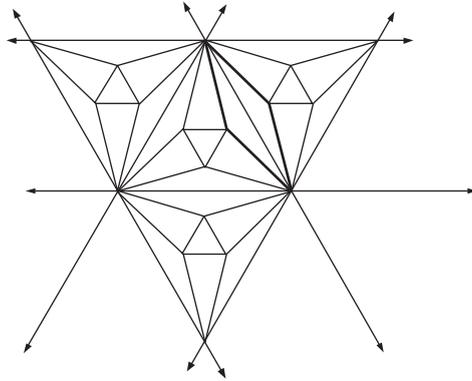


Fig. 5. The *flower graph* is obtained from tiling the plane with this pattern.

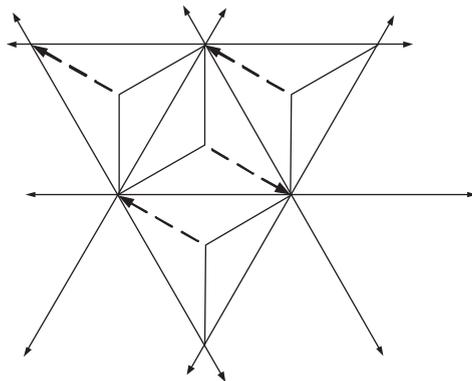


Fig. 6. The *flower graph* requires a non-symmetric embedding.

2.2. On the existence of 3-connected planar spanning subgraphs

Assuming the conjecture is true, we would then naturally ask which 3-connected graphs have a 3-connected planar spanning subgraph (and hence have a greedy embedding by the conjecture). Are there any 3-connected graphs with no 3-connected planar spanning subgraph?

Such graphs exist, and $K_{3,3}$ is the simplest example: it is both minimally non-planar and minimally 3-connected. However, the following result shows that essentially this is the only counterexample:

Theorem 2. *If a 3-connected graph does not have a $K_{3,3}$ minor, then it has a 3-connected planar spanning subgraph.*

Proof. Suppose the graph has no 3-connected planar spanning subgraph, and consider a maximal planar spanning subgraph G ; it must have at least 5 vertices and a cut set of two vertices, a and b . Adding back an edge e to G results in a non-planar graph. By Kuratowski's theorem [7] and the hypothesis, $G + e$ has a K_5 minor. Consider these five vertices; at least three of them are neither a nor b . There are three cases, depending on the distribution of these three vertices on the connected components of $G - \{a, b\}$. If they are all on the same component, then G has a K_5 minor and is thus non-planar, contradicting our hypothesis that G is planar. And if they are on two or three different components, they cannot be fully connected to each other via the single edge e . \square

3. Polyhedral routing

Is the conjecture true if we allow embedding in higher dimensions? We do not know. But we show here that any 3-connected planar graph (or any graph containing such) can be embedded in three dimensions so that greedy routing works, albeit by a different interpretation of the “distance” $d(u, v)$: it is now not the Euclidean distance of the images of u and v , but *the negative dot product of the coordinates of those images*.

Theorem 3. *Any graph containing a 3-connected planar graph has a greedy embedding e in \mathfrak{R}^3 , provided that we define $d(u, v) = -\mathbf{e}(u) \cdot \mathbf{e}(v)$.*

Proof. Steinitz showed in 1922 that every 3-connected planar graph is the edge graph of a three-dimensional convex polytope [16]. Such representations are by no means unique, and further work has shown that such a polytope exists even under the constraint that all edges must be tangent to a sphere [14]. Now, a crucial property of a convex polytope is that each vertex of the polytope has a supporting (“tangent”) hyperplane that contains that vertex but which does not otherwise intersect the polytope. Furthermore, if the polytope has all edges tangent to a sphere centered at the origin (as the Koebe–Andre’ev representation does), then the three-dimensional coordinates of the vertices of the polytope also serve as the normal vector of the supporting hyperplane, as illustrated in Fig. 7.

It remains to show that, for any two distinct vertices s and t of such a polytope, there is a neighbor v of s such that $d(v, t) < d(s, t)$. However, this follows immediately from convexity: consider the linear function $L_t(x)$ defined by the hyperplane normal to the vector $\mathbf{e}(t)$. If no such vertex v exists, then $\mathbf{e}(s)$ is a local optimum of $L_t(x)$, and, by convexity, the global optimum. However, since $L_t(x)$ is defined by a hyperplane normal to $\mathbf{e}(t)$, $\mathbf{e}(t)$ is the unique optimum of $L_t(x)$, and thus $s = t$, contradiction. \square

3.1. Experiments

Employing the above embedding in a practical algorithm is confounded by two factors. First, in practical settings, the connection graph need not necessarily be 3-connected or planar. Second, the construction of the Koebe–Andre’ev embedding is not easily obtainable in a distributed fashion [9], nor is it obvious how such a scheme can be made robust to the intermittent node failures common in ad hoc networks. For these reasons, the significance of Theorem 3 remains largely theoretical (Fig. 8).

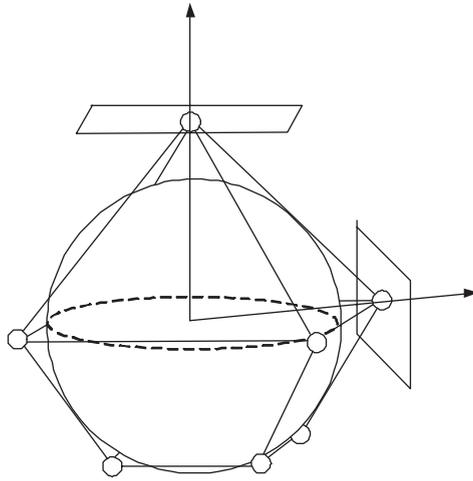


Fig. 7. Koebe–Andre'ev Steinitz representation with supporting hyperplanes.

Radius (r)	Original	Euclidean	Spherical
0.27	2.347	1.665	1.368
0.28	1.932	1.526	1.253
0.3	1.508	1.486	1.240
0.4	0.165	0.764	0.683
0.5	0.004	0.468	0.469

Fig. 8. The average percentage of paths that are *not* successfully routed.

As an incremental step toward proving the practicality of the spherical embedding, we ran experiments on a variant of the *rubber band* algorithm [12] in which the boundary nodes are evenly spaced on the equator of a unit sphere, and the positions of the remaining nodes are determined by a physical simulation assuming that all edges are equally strong rubber bands that are stretched along the surface of the sphere. For routing, we use the dot-product distance of Theorem 3 rather than Euclidean distance. We surmised that when the graph was 3-connected and nearly planar, the resulting embedding would be similar to a Koebe–Andre'ev embedding projected onto a hemisphere.

For our simulation, we began with a random 3-connected geometric graph [11] G of 100 nodes placed within a unit square region and with radius of transmission r . We chose three random nodes as boundary nodes (note these need not be on a single face even if G were planar) and then ran both the Euclidean rubber band algorithm in two dimensions, and

the spherical rubber band algorithm in three dimensions.² For varying choices of r , we determine the average percentage (over 100 trials) of paths that run into a void and do not terminate at the correct destination.

The results in Fig. 2 show that the spherical version of the embedding is a slight improvement over the classical rubber band algorithm under various radii of transmission (increasing r increases the connectivity of the graph). For small values of r , the virtual coordinates of the Euclidean and spherical embeddings perform better than the actual geographic coordinates, an effect which was also noted by Rao et al. [12]. For larger values of r , the graph has larger cliques and becomes less planar. This results in decreased performance of the embeddings as r is increased, and decreased disparity in the performance of the two embeddings.

4. Face routing

When greedy routing fails, most existing geographic routing algorithms resort to *face routing*: they circumnavigate a face (rather, what appears to the protocol to be a face) until greedy routing can resume [4,6,12]; such algorithms are either quite complicated, or are not guaranteed to work. In this section we point out that a convex embedding of a 3-connected planar graph yields a very simple and rigorous form of face routing.

Given a 3-connected planar graph G with a convex embedding, we let each vertex u record its own position, the position of each neighbor v , as well as the names of the two faces incident upon the edge uv . In the course of routing, we proceed towards destination t by greedy routing. If a node $u \neq t$ is reached without a neighbor that is closer to t , a *face routing* phase is initiated, as follows: by the location of itself, of t , and of its neighbors, u identifies the two clockwise consecutive neighbors v, w such that t lies in the angle formed by vuw , and the face f where the ray ut lies (it is the face incident upon both edges uv and uw). It then circumnavigates the face f , say clockwise, until one of the following two events occurs:

- A vertex u' is reached with $d(u', t) < d(u, t)$, where by u we mean the node at which face routing started; in this case greedy routing resumes.
- A vertex a is reached whose edge ab on face f is closer to t than both $d(a, t)$ and $d(b, t)$. Face routing then is continued in the other face incident upon ab .

To see that one of these events will occur, recall that the distance from t to points of a convex polygon (the face f), not including t in its interior, is achieved either at a vertex or at an edge of the polygon. Note that this algorithm bears some resemblance to compass routing [4], though it is much simpler due to our assumption of a convex embedding (Fig. 9).

Theorem 4. *Given a 3-connected planar graph and a convex embedding of that graph, the above routing algorithm always terminates at t .*

² The random choice of boundary nodes differs from the more complicated scheme of Rao et al. [12], however, we wished to compare the performance of the two embedding schemes with a version of the algorithm that would be simplest to implement in practice.

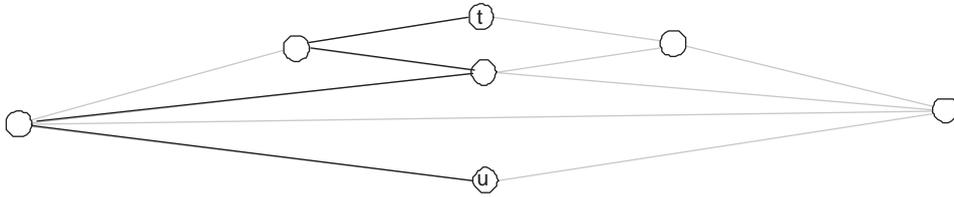


Fig. 9. Face routing on a convex embedding.

Proof. Define the distance from a face f to t , $d(f, t)$ to be the minimum of $|ta|$ over all points a in f . If the minimum occurs along an edge ab of f , we say that ab realizes $d(f, t)$.

We must show that during face routing, no face is circumnavigated more than once. At each circumnavigation of a face f , either greedy routing resumes, in which case there is nothing to prove, or we arrive at an edge that realizes $d(f, t)$. When face routing is initiated at u , we know that $d(f, t) < d(u, t)$ (because u is a vertex of f and the ray ut lies within f). Furthermore, when face routing is continued from face f to face f' , again $d(f', t) < d(f, t)$ (because this happens when the common edge of f and f' realizes $d(f, t)$). It follows that the distance to t (in this expanded version also comprising faces) monotonically decreases during the algorithm. \square

5. Discussion

We have thus far been unable to prove the conjecture, even for the special case when the graph is a maximally planar graph (all faces are triangles). One particularly attractive property of 3-connected planar graphs is a lemma of Thomasson stating that all 3-connected planar graphs of more than 4 vertices have an edge that can be contracted to yield another 3-connected planar graph. We initially wished to inductively construct the embedding based on this lemma (and some of its variants). However, all efforts in this direction have failed. While we have confidence in the truth of the conjecture, we suspect that its proof is non-trivial.

While the conjecture proposes a sufficient condition for a greedy embedding, it is by no means necessary. There are graphs with no 3-connected planar subgraph for which we can produce a greedy embedding ($K_{3,3}$). Recall also that any graph with a Hamiltonian path has a greedy embedding on a line. Properly characterizing graphs with a greedy embedding therefore remains very much an open question, and we view the results of this paper only as a starting point.

Acknowledgements

We would like to thank Janos Pach, Nati Linial, and Vladlen Koltun for their ideas and insights regarding the conjecture. We also thank the anonymous reviewers for their helpful comments.

References

- [1] P. Bose, P. Morin, I. Stojmenovic, J. Urrutia, Routing with guaranteed delivery in ad hoc wireless networks, *Wireless Networks* 7 (6) (2001) 609–616.
- [2] D. Estrin, R. Govindan, J.S. Heidemann, S. Kumar, Next century challenges: scalable coordination in sensor networks, *Mobile Comput. Networking* (1999) 263–270.
- [3] B. Karp, H.T. Kung, GPSR: greedy perimeter stateless routing for wireless networks, *Mobile Comput. Networking* (2000) 243–254.
- [4] E. Kranakis, H. Singh, J. Urrutia, Compass routing on geometric networks, in: *Proc. 11th Canad. Conf. on Computational Geometry*, Vancouver, 1999, pp. 51–54.
- [5] F. Kuhn, T. Moscibroda, R. Wattenhofer, Unit disk graph approximation, in: *Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M)*, 2004.
- [6] F. Kuhn, R. Wattenhofer, Y. Zhang, A. Zollinger, Geometric ad-hoc routing: of theory and practice, in: *Proc. 22nd ACM Int Symp. on the Principles of Distributed Computing, (PODC)*, 2003.
- [7] K. Kuratowski, Sur le probl'eme des courbes gauches en topologie, *Fund. Math.* 15 (1930) 271–283.
- [8] N. Linial, L. Lovasz, A. Wigderson, Rubber bands, convex embeddings and graph connectivity, *Combinatorica* 8 (1988) 91–102.
- [9] L. Lovasz, Steinitz representations of polyhedra and the colin de verdere number, *J. Combin. Theory B* (82) (2001) 223–236.
- [10] T. Moscibroda, R. O'Dell, M. Wattenhofer, R. Wattenhofer, Virtual coordinates for ad hoc and sensor networks, in: *Proc. of 2nd Joint Workshop on Foundations of Mobile Computing (DIALM-POMC)*, 2004.
- [11] M. Penrose, *Random Geometric Graphs*, Oxford University Press, Oxford, 2003.
- [12] A. Rao, C. Papadimitriou, S. Shenker, I. Stoica, Geographic routing without location information, in: *Proc. of the 9th Annual Internat. Conf. on Mobile Computing and Networking*, ACM Press, New York, 2003, pp. 96–108.
- [13] E. Royer, C. Toh, A review of current routing protocols for ad-hoc mobile wireless networks, *IEEE Personal Commun.* (1999).
- [14] W. Thurston, *Three-dimensional Geometry and Topology*, Princeton Mathematical Series, Vol. 35, Princeton University Press, Princeton, NJ, 1997.
- [15] W.T. Tutte, Convex representations of graphs, *Proc. London Math. Soc.* 10 (1960) 304–320.
- [16] G.M. Ziegler, *Lectures on Polytopes*, Springer, Berlin, 1995.