On a Conjecture Related to Geometric Routing

Christos H. Papadimitriou* and David Ratajczak

Computer Science Division
University of California at Berkeley
Berkeley, CA 94720-1776, USA,
{christos,dratajcz}@cs.berkeley.edu

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 subgraph, 2-dimensional virtual coordinates for the nodes can be found for which greedy geographic routing is guaranteed to work. We discuss this conjecture and its equivalent forms. We show a weaker result, namely that for any network containing a 3-connected planar subgraph, 3-dimensional virtual coordinates always exist enabling a form of greedy routing inspired by the simplex method; we provide experimental evidence that this scheme is quite effective in practice. We also propose a rigorous form of face routing based on the Koebe-Andre'ev-Thurston theorem. Finally, we show a result delimiting the applicability of our approach: any 3-connected $K_{3,3}$-free graph has a planar 3-connected subgraph.

1 Introduction

Routing in the Internet is typically based on IP addresses. In ad hoc networks, however, there is no such universally known system of addresses. *Geometric routing* [1,2] is a family of routing algorithms using the geographic coordinates of the nodes as addresses for the purpose of routing. One such algorithm is *greedy routing* which is attractive for its simplicity; each node gives the packet to the neighbor that has the closest Euclidean distance to the destination. Unfortunately, purely greedy routing sometimes fails to deliver a packet because of the phenomenon of "voids" (nodes with no neighbor closer to the destination). To recover from failures of greedy routing, various forms of *face routing* have been proposed [2], in which the presence of a void triggers a special routing mode until the greedy mode can be reestablished. For example, a recent paper [1] proposes a variant of this idea that is guaranteed to deliver the message, and in fact 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 future ad hoc networks can rely on the availability of precise geographic

* Supported by NSF ITR grant CCR-0121555.

S. Nikoletseas and J. Rolim (Eds.): ALGOSENSORS 2004, LNCS 3121, pp. 9–17, 2004.

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. [3] 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 [4] and used often in graph theory [5]. It was noted, on the basis of extensive experimentation, that this approach makes greedy routing much more reliable (in section 3 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 [3] in geometric graph theory, no theoretical results and guarantees are known for such schemes. *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.*

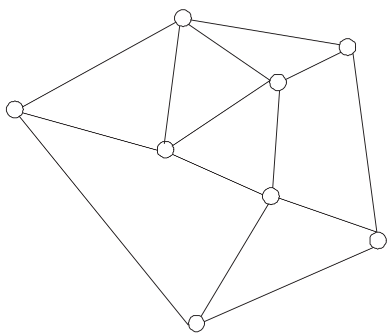

Fig. 1. A *greedy* embedding.

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 Figure 1. It has the following property: given any nodes s and t, there is a neighbor of s that is closer (in the sense of 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 [4] (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 adding edges only improves the embeddability of a graph, the conjecture extends immediately to any graph with a 3-connected planar subgraph.

2 The Conjecture

Let G be a graph with nodes embedded in $\Re^k$ and let d denote the k-dimensional Euclidean distance. We say that 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 $i = 1, \ldots, m$. We propose the following:

Conjecture 1 (Weak). 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 from s to t.

We call such an embedding a *greedy embedding*, for reasons which are clear from the following theorem.

Theorem 1. *The following are equivalent forms of conjecture 1:*

1. *Any 3-connected planar graph has an embedding on the plane in which for any two nodes s and t there is a neighbor r of s such that $d(r, t) < d(s, t)$.*
2. *Any 3-connected planar graph has an embedding on the plane in which greedy routing will successfully route packets between any source and destination.*
3. *Any graph that contains a 3-connected planar graph has a greedy embedding.*

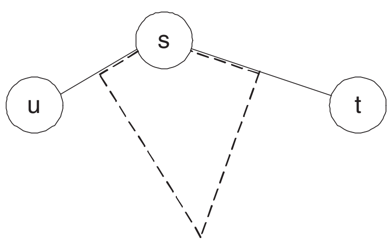

Fig. 2. The forbidden region.

A *convex embedding* of a planar graph is a planar embedding in which all faces, including the external face, are convex. It is known that every 3-connected planar graph has a convex embedding [4] and we additionally conjecture that:

Conjecture 2 (Strong). All 3-connected planar graphs have a greedy convex embedding.

We immediately have the following theorem which is a useful shortcut for testing if an embedding is greedy:

Theorem 2. *A convex embedding of a 3-connected planar graph is greedy if and only if for any obtuse angle about a face (formed by nodes u, s and t), the intersection of the half-planes defined by su (containing t), st (containing u), and the perpendicular bisectors of the segments su and sv (containing s) contain no other vertex of the graph.*

This is illustrated in Figure 2.

Some Counterexamples

We have obtained a family of counterexamples based on a simple lemma.

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

Proof. Otherwise, u has no neighbor that is closer to s than itself.

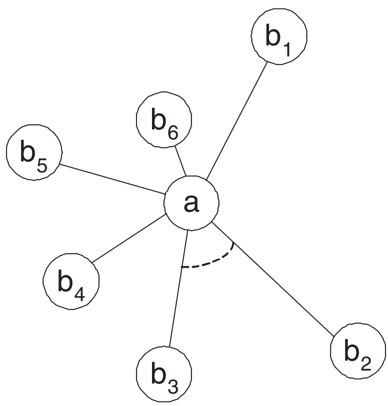

Fig. 3. $K_{1,6}$ has no greedy embedding.

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 $b_1, \ldots, b_6 \in B$ 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 Figure 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.

These counterexamples imply that the hypotheses of the conjecture are necessary, in that there exist counterexamples that are planar but not 3-connected ($K_{2,11}$), or 3-connected but not planar ($K_{3,16}$); also, they show that high-connectivity alone does not guarantee a greedy embedding.

In our quest to obtain 3-connected planar counterexamples, we have investigated other types of graphs, such as the graph obtained by tiling the plane with

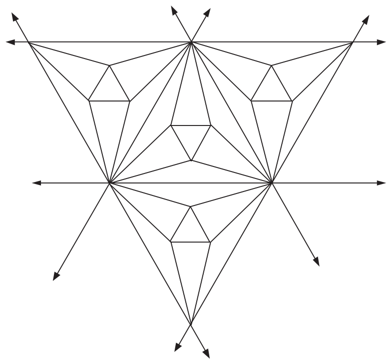

Fig. 4. A graph requiring a non-symmetric embedding.

K_4, or the graph shown in Figure 4. Embedding these graphs require breaking their symmetry and resorting to non-local arguments. However, they do indeed have convex greedy embeddings.

On the Existence of 3-Connected Planar Subgraphs

Assuming the conjecture is true, we may ask which 3-connected graphs have a 3-connected planar subgraph (and hence have a greedy embedding by the conjecture). Are there any 3-connected graphs with no 3-connected planar 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 3. *If a 3-connected graph does not have a $K_{3,3}$ minor, then it has a 3-connected planar subgraph.*

Proof. Suppose the graph has no 3-connected planar subgraph, and consider a maximal planar 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 [] 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 nonplanar, contradiction. And if they are on two or three different components, they cannot be fully connected to each other via the single edge e.

3 Polyhedral Routing

In this section we consider a "greedy" form of routing on 3-connected planar graphs, where the coordinates are 3-dimensional, and the distance between two points corresponds to their dot product.

Steinitz showed in 1922 that every 3-connected planar graph is the edge graph of a 3-dimensional convex polytope [6]. 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 circle [7]. A crucial property of a convex polytope, employed in the simplex method, is that each vertex of the skeleton has a supporting hyperplane that is tangent to that vertex but which does not intersect the polytope. Furthermore, every other vertex has a neighbor which is closer (in perpendicular distance) to that hyperplane. If the polytope has all edges tangent to a circle centered at the origin (as the Koebe-Andre'ev representation does), then the 3-dimensional coordinates of the corners of the polytope also serve as the normal vector of the supporting hyperplane, as illustrated in Figure 5.

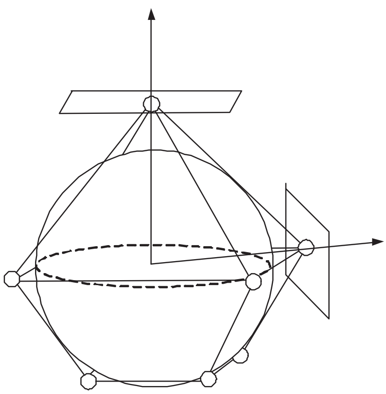

Fig. 5. Koebe-Andre'ev Steinitz representation with supporting hyperplanes.

This implies the following greedy algorithm. Given a 3-connected planar graph G, calculate a Koebe-Andre'ev embedding in 3-dimensions about a circle centered at the origin. For a node v, let $p(v)$ denote its 3-dimensional coordinate in the embedding. When routing to t, s will forward a packet to the neighbor v such that $p(v) \cdot p(t)$ is maximized.

By the convexity of the polytope and the fact that $p(v)$ defines a supporting hyperplane for v, there will always be a neighbor v of s such that $p(v) \cdot p(t) > p(s) \cdot p(t)$. Since $p(v) \cdot p(t)$ is maximized when $v = t$, the above routing algorithm must always make progress until it terminates at the desired destination.

Experiments

The implementation of the above algorithm in a practical setting 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 [8].

Instead, we consider a variant of the *rubber band* algorithm [3] 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 rather than Euclidean distance.

For our simulation, we began with a random 3-connected geometric graph [9] 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 2-dimensions, and the spherical rubber band algorithm in 3-dimensions. 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 of Figure 3 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. [3]. 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 peformance of the two embeddings.

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. 6. The average percentage of paths which are *not* successfully routed.

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 [1,3]; such algorithms are either quite complicated, or are not guaranteed to work. In this section we point out that, using a recent strong planar embedding result, we can show that any 3-connected planar graph can be embedded so that a simple and rigorous form of face routing works.

The following statement is equivalent to the existence of the Koebe-Andre'ev Steinitz representation, but restated in a simplified way for our purposes:

Theorem 4. (Koebe, Andre'ev, Thurston) [7]: *Any 3-connected planar graph and its dual can be simultaneously embedded on the plane so that each face is a convex polygon with an inscribed circle whose center coincides with the vertex of the dual corresponding to the face, and so that edges are perpendicular to their dual edges.*

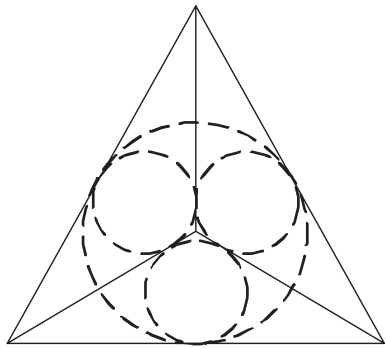

Fig. 7. Koebe-Andre'ev planar representation with inscribed circles.

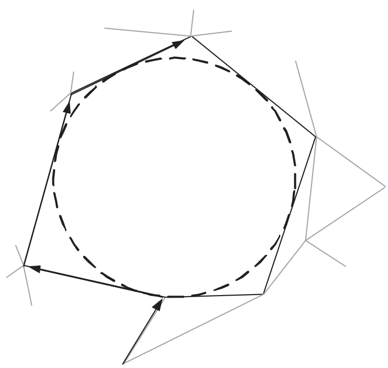

Fig. 8. Face routing on a Koebe-Andre'ev embedding.

Figure 7 illustrates an example of this embedding. This result gives rise to the following rigorous face routing protocol: each vertex knows, besides the coordinates of its neighbors, the coordinates of all dual vertices representing the faces adjacent to it; each such face points to the neighbor vertex that comes next in clockwise order. When greedy routing fails, then the vertex employs face routing: it selects the face (dual vertex) that is closest to the destination, and sends the packet clockwise around that face. When a vertex is found with a neighbor that is closer to the destination than the vertex that began the face routing, greedy routing is re-established. This process is illustrated in Figure 8. We omit the proof of the following:

Theorem 5. *Any graph containing a 3-connected planar graph can be embedded on the plane so that the above face routing protocol always succeeds.*

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.

Acknowledgments. We would like to thank Janos Pach and Nati Linial for their ideas and insights regarding the conjecture.

References

1. Kuhn, F., Wattenhofer, R., Zhang, Y., Zollinger, A.: Geometric Ad-Hoc Routing: Of Theory and Practice". In: Proc. 22^{nd} ACM Int. Symposium on the Principles of Distributed Computing (PODC). (2003)
2. Karp, B., Kung, H.T.: GPSR: greedy perimeter stateless routing for wireless networks. In: Mobile Computing and Networking. (2000) 243–254
3. Rao, A., Papadimitriou, C., Shenker, S., Stoica, I.: Geographic routing without location information. In: Proceedings of the 9th annual international conference on Mobile computing and networking, ACM Press (2003) 96–108
4. Tutte, W.T.: Convex representations of graphs. Proceedings London Math. Society **10** (1960) 304–320
5. Linial, N., Lovasz, L., Wigderson, A.: Rubber bands, convex embeddings and graph connectivity. Combinatorica **8** (1988) 91–102
6. Ziegler, G.M.: Lectures on Polytopes. Springer-Verlag, Berlin (1995)
7. Thurston, W.: Three-dimensional Geometry and Topology. Number 35 in Princeton Mathematical Series. Princeton University Press, Princeton, NJ (1997)
8. Lovasz, L.: Steinitz representations of polyhedra and the colin de verdiere number. Journal of Combinatorial Theory **B** (2001) 223–236
9. Penrose, M.: Random Geometric Graphs. Oxford University Press (2003)