

On the Pitfalls of Geographic Face Routing

Young-Jin Kim[†] Ramesh Govindan[†]
[†]University of Southern California
Los Angeles, CA 90089
{youngjki, ramesh}@usc.edu

Brad Karp*
*Intel Research/CMU
Pittsburgh, PA 15213
bkarp@cs.cmu.edu

Scott Shenker[‡]
[‡]UCB/ICSI
Berkeley, CA 94704
shenker@icsi.berkeley.edu

ABSTRACT

Geographic face routing algorithms have been widely studied in the literature [1, 8, 13]. All face routing algorithms rely on two primitives: *planarization* and *face traversal*. The former computes a planar subgraph of the underlying wireless connectivity graph, while the latter defines a consistent forwarding mechanism for routing around “voids.” These primitives are known to be provably correct under the idealized unit-disk graph assumption, where nodes are assumed to be connected if and only if they are within a certain distance from each other.

In this paper we classify the ways in which existing planarization techniques fail with realistic, non-ideal radios. We also demonstrate the consequences of these pathologies on reachability between node pairs in a real wireless testbed. We then examine the various face traversal rules described in the literature, and identify those [12, 16] that are robust to violations of the unit-disk graph assumption.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols - Routing Protocols

General Terms

Algorithms, Reliability, Performance, Experimentation

Keywords

Geographic routing, Planarization, Face changes, Cross-Link Detection

1. INTRODUCTION

Geographic routing algorithms for wireless ad hoc networks have been shown to scale better than other alternatives: they require per-node state that depends only on network density and not on network size. There is a broad literature on geographic routing algorithms: from initial sketches suggesting routing using position information [2, 10]; to the first practical, detailed proposals, including

GFG [1], GPSR [8], and the GOAFR+ family of algorithms [13]; to refinements of these proposals for efficiency [4], robustness under real network conditions [15, 22], and even routing geographically when node location information is unavailable [18, 19].

In many of these algorithms, nodes forward packets to the neighbor closest to the destination. This greedy forwarding fails in the presence of radio *voids*. These algorithms recover from such failures by successively traversing the faces of the planar subgraph of the underlying connectivity graph. To compute the planar subgraph, these algorithms use graph planarization algorithms that are amenable to distributed implementation. These planarization algorithms [3, 4, 23] rely purely on node location information to determine whether or not links to neighbors belong in the planar subgraph.

A common assumption underlying the planarization and face traversal algorithms is that connectivity between nodes can be described by unit graphs.¹ In such graphs, a node is always connected to all nodes within its fixed, “nominal” radio range, and is never connected to nodes outside this range. However, in wireless networks where geographic routing is likely to be deployed, the unit-disk graph assumption can be violated. For example, geographic routing has been proposed for use as a routing primitive and as a building block for data storage and flexible query processing in sensor networks [17, 20]. In these networks, radio-opaque obstacles and multi-pathing cause non-uniform radio ranges. Furthermore, the ad hoc localization techniques proposed for these networks can result in imprecise node location, which in turn violates the unit-disk graph assumption [22].

Planarization and face traversal failures, as summarized in Table 1, can result from violations of the unit-disk graph assumption. These failures can then result in persistent routing failures, where geographic routing fails to find a path for at least one source-destination pair. In this paper, we first taxonomize the causes and effects of planarization failures. We then validate, using an actual implementation of GPSR [8], a previously proposed geographic face routing protocol, that such failures arise in real wireless networks. In our experiments, GPSR left more than 30% of node pairs disconnected. The “mutual-witness” extension to GPSR still did not offer complete reachability, leaving 10% of node pairs disconnected. These results motivate CLDP [9], a qualitatively different approach to eliminating crossing edges from the network graph.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DIALM-POMC'05, September 2, 2005, Cologne, Germany.
Copyright 2005 ACM 1-59593-092-2/05/0009 ...\$5.00.

¹Some recent literature [15] has relaxed this assumption to allow for “quasi” unit disk graphs, which always exhibit connectivity within a short radius r ; exhibit probabilistic connectivity within an enclosing ring-shaped region; and exhibit no connectivity beyond this ring. The technique only is scalable when $r \geq 1/\sqrt{2}$; for lesser values of r , virtual links may be comprised of increasingly long paths of physical hops.

Planarization (Section 3.1)	unidirectional links in subgraphs disconnected links in subgraphs cross links in subgraphs collinear links in subgraphs
Face traversal (Section 4.1)	Incorrect face changes (GPSR, GFG, GOAFR+)

Table 1: Classification of pathologies

We next examine the robustness of face traversal algorithms to violations of the unit-disk graph assumption. We first find that the commonly used “right-hand rule” can fail in the presence of collinear links, then propose and prove the correctness of a revised right-hand rule. More importantly, we extensively examine the rule used to traverse successive faces of the planar subgraph (called the *face-change rule*). We find that only two of the face-change rules proposed in the literature are correct in realistic settings.

This work, together with [9], establishes that for geographic routing to work in realistic wireless settings, several of its key elements bear careful re-examination and redesign.

2. PRELIMINARIES

We now briefly review the principles underlying the class of geographic routing algorithms that use face traversal to circumnavigate voids. Examples of protocols in this family include GFG [1], GPSR [8], and GOAFR+ [13]. We refer to this family of algorithms simply as *geographic routing*.² In general, this family of algorithms relies on three distinct components: *greedy forwarding*, *face traversal*, and *planarization*. The following paragraphs describe these components. In what follows, we assume that each node acquires its own position using GPS devices or using ad hoc localization techniques [21], and learns of its neighbors’ positions using a simple beaconing protocol.

Packets stamped with the location of a destination are forwarded using a *greedy traversal* when possible. In this greedy traversal, a node sends a packet to its neighbor that is geographically closest to the destination. If there is no neighbor that is closer to the destination, the packet has reached a local minimum, and greedy traversal fails. *Face traversal* is then invoked to recover from local minima. When face traversal (described below) visits a node closer to the destination than the local minimum, it falls back to greedy traversal. We call the node at which this happens the *first closure*.

Face traversal works correctly on a network graph that has no crossing links, *i.e.*, a *planar* graph. A planar graph consists of *faces*, enclosed polygonal regions bounded by links. Face traversal uses two primitives to traverse a planar graph: the *right-hand rule* and *face changes*. Figure 1 describes the right-hand rule, which dictates that upon receiving a packet on a link, the receiving node forward that packet on the first link it finds after sweeping counter-clockwise about itself from the ingress link. Now consider the planar graph in Figure 2, in which the source node S and destination node D are as indicated. Observe that the line segment \overline{SD} must cut a series of faces in the planar graph. Face traversal intends to successively walk the faces cut by this line. The numbering of faces in Figure 2 shows the order in which faces are traversed from S to D on that planar graph. To ensure correct traversal of faces, the right-hand rule must be augmented with a rule that describes when to switch to an adjacent face. There are several rules that have been proposed in the literature, and we review them in Section 4.

²We note that there exist other routing algorithms that make use of position information, such as LAR [11], but we restrict our view to algorithms in which a node forwards to a single neighbor on the basis of geographic information.

Finally, the *planarization* component eliminates crossing links to generate the planar sub-graph required for face traversal. Geographic routing algorithms planarize network graphs by using the Gabriel Graph (GG) [3], the Relative Neighborhood Graph (RNG) [23], or Restricted Delaunay Graph (RDG) [4]. These graph constructs provably yield a connected, planar graph so long as the connectivity between nodes obeys the *unit-disk assumption*. Here is how they work. Consider a link in the network graph between two nodes A and B in Figure 3. Both A and B must decide whether to eliminate the link between them in the planar graph or not. For both the GG and RNG, node A searches its 1-hop neighbor list for any *witness* node W that lies within a particular geometric region. If one or more witnesses are found, the link (A, B) is eliminated in the planar graph. For the GG, the region where a witness must exist to eliminate the link is the circle whose diameter is line segment \overline{AB} . For the RNG, this region is the *lune* defined by the intersection of the two circles centered at A and B , each with radius $|\overline{AB}|$.

As already described, currently known geographic routing algorithms have relied on the unit-graph assumption. Hence, whether wireless network graphs in reality conform to the assumption is a question of great importance. Several experimental studies on wireless ad hoc networks [24,25] have shown that real radio ranges are irregular—that is, the unit-graph assumption is violated. More recent research [26] has explored causes of radio range irregularity.

3. PLANARIZATION IN REAL ENVIRONMENTS

The planarization techniques used by geographic face routing protocols are guaranteed to produce a connected planar subgraph when the connectivity graph is a unit-disk graph. In this section, we taxonomize the failures in existing planarization techniques when the unit-disk assumption is violated. We briefly discuss a qualitatively different planarization technique (published elsewhere [9]) that does not rely on the unit-disk assumption. Finally, we quantify the impact of planarization failures on routing efficacy in two real-world deployments of geographic routing protocols.

3.1 Planarization Failures

The unit-disk assumption can, in practice, be violated in one of two ways. A node can have an incorrect estimate of its position. Alternatively, its radio communication range may be irregular as a result of radio-opaque obstacles or transceiver differences. These deviations from ideal behavior can cause three kinds of failures in the planarization algorithms described above: they can leave unidirectional links in the GG/RNG-derived subgraph, incorrectly remove links from this subgraph, or leave cross-links in the subgraph. In what follows, we illustrate failures as a result of applying the GG algorithm. We have observed each type of failure in an experimental setting, as we describe later in the section.

To our knowledge, this is the most complete taxonomy of planarization failures to date. Prior work [22] has described the latter two failures resulting from localization errors.

Unidirectional links

Consider the upper left subfigure of Figure 4. The distance between C and A is less than that between C and B . C and A cannot communicate with each other due to a radio-opaque obstacle between them. When node C receives beacons from node B , it executes the GG planarization (Section 2). But node C does not know of the existence of node A , and it keeps link (C, B) in its GG-derived subgraph. Furthermore, B notices that A is a witness, and removes link (B, C) from its subgraph. Thus, the GG-derived subgraph is left

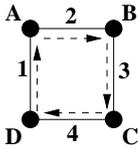


Figure 1: Right-hand rule. A sweeps counterclockwise from link 1 to find link 2, forwards to B , &c.

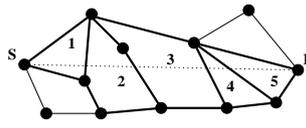


Figure 2: The faces progressively closer from S to D along line segment \overline{SD} , numbered in the order visited. Faces cut by \overline{SD} are bordered in bold.

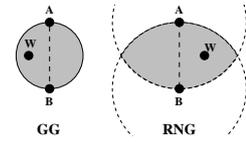


Figure 3: Geometric definitions of the GG and RNG. A witness must fall within the shaded circle (GG) or the shaded lune (RNG) for edge (A,B) to be eliminated in the planar graph.

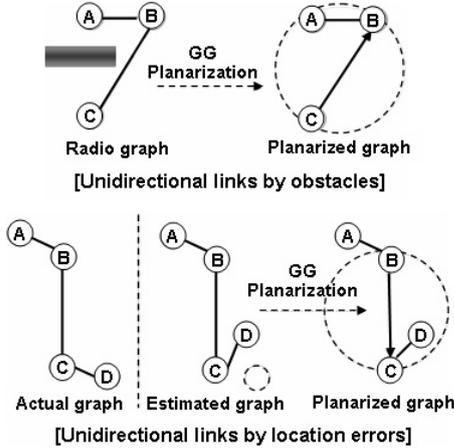


Figure 4: Unidirectional links in subgraphs.

with a unidirectional link, denoted by an arrow in the upper-right subfigure of Figure 4.

The GG-derived subgraph can contain the same artifact when nodes have imprecise location information. The lower left subfigure of Figure 4 shows the actual location of 4 nodes, as well as their “estimated” topology (based on each node’s estimate of its position). In the estimated topology, D has incorrect position information. C assumes D to be a witness node for its link to B , and does not add (C,B) to the GG-derived subgraph. By contrast, B cannot communicate with D , and thus preserves link (B,C) in the GG-derived subgraph. As a result, the GG-derived subgraph has a unidirectional link as shown in the lower right subfigure of Figure 4.

These unidirectional links can produce an infinite loop during face traversal, as shown in Figure 7.

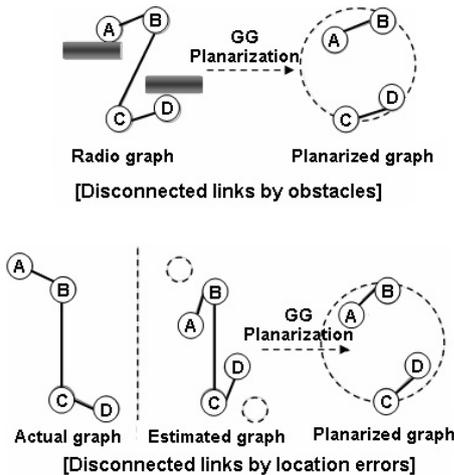


Figure 5: Disconnected links in subgraphs.

Disconnected links

Radio-opaque obstacles can cause a link to be incorrectly removed from the GG-derived subgraph. In the upper part of Figure 5, obstacles prevent B and D , as well as A and C , from communicating with each other. C realizes that link (C,B) is witnessed by node D and removes that link from the GG-derived subgraph. B realizes that link (B,C) is witnessed by node A and removes that link from the GG-derived subgraph. This results in a disconnected subgraph, as shown in the upper right of Figure 5.

Incorrect node location can also produce a similar artifact. The lower left subfigure of Figure 5 illustrates both the actual locations and the estimated topology of four nodes. In the estimated topology, A and D have incorrect position information. As with the previous example, the GG-derived subgraph is disconnected, as shown in the lower-right subfigure of Figure 5.

A disconnected GG-derived subgraph can cause face traversal failures, as shown in Figure 8. In this figure, there exists a path from S to D . Greedy traversal towards D encounters a local minimum. However, face traversal cannot reach D , since the GG-derived subgraph is partitioned.

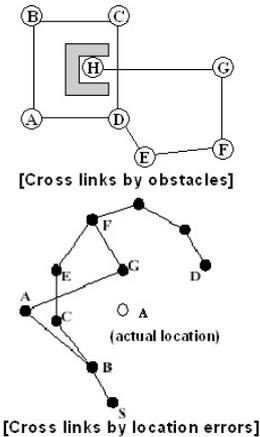


Figure 6: Cross links in subgraphs.

Cross links in subgraphs

Finally, in the presence of non-ideal radio behaviors or incorrect location information, the GG and RNG algorithms can leave cross-links in the derived subgraphs. In the upper part of Figure 6, H has only a link to G , as a result of radio-opaque obstacles. Even though the link (H,G) intersects link (C,D) , the GG planarization fails to remove the crossing links, as C and D cannot see H , and H and G cannot see nodes C and D . Similarly, in the lower part of Figure 6, A is actually positioned below G (white circle). However, A incorrectly assumes that its own position is near C and E . The GG planarization leaves the crossing links (G,A) and (C,E) .

Crossing links in derived subgraphs can cause persistent routing failures, as illustrated in Figure 9.

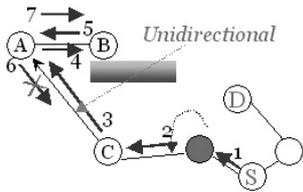


Figure 7: Routing failure caused by unidirectional links.

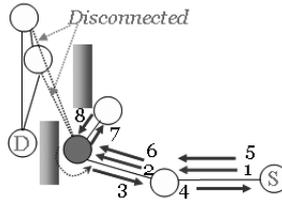


Figure 8: Routing failure caused by disconnected links.

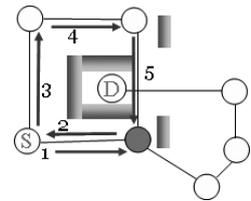


Figure 9: Routing failure caused by cross links.

3.2 The Mutual Witness Procedure

Some of the planarization failures arise because two ends of a link have inconsistent information about witnesses. The *mutual witness protocol* [6, 7, 22] attempts to rectify this. When node A considers whether to keep a link (A, B) in the subgraph, mutual witness dictates that A only eliminate link (A, B) if there exists at least one witness in the RNG or GG region that is visible *both* to A and B . This fact may be directly verified with local communication: if all nodes broadcast their neighbor lists (only a single hop), then all nodes may verify whether a particular neighbor shares a particular other neighbor.

As our examples below show, the mutual witness protocol is not perfect; on non unit-disk graphs, it converts *unidirectional links* or *disconnected links* into *cross links* in the subgraph, leaves *some cross links* in the subgraph, and can introduce collinear links. Consider the first row of Figure 10. The GG planarization without mutual witness results in a unidirectional link (A, C) . If mutual witness is applied instead, node C keeps link (C, A) because A does not also observe the witnesses B and C . This can leave a cross-link in subgraph (the right-most of the first row). In the second row of Figure 10, the GG planarization would remove links (A, D) and (B, C) in the subgraph (the middle subfigure of the second row), but if mutual witness is applied to the subgraph, the crossing links (A, D) and (B, C) are left in the derived subgraph. Finally, the third row of Figure 10 shows an example where both the GG planarization and mutual witness leave crossing links in the derived subgraph.

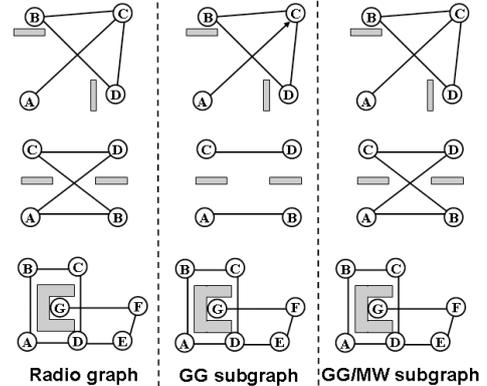


Figure 10: Mutual witness is not enough.

3.3 Experimental Validation of Planarization Failures

We implemented GPSR in TinyOS [5], the event-driven operating system used on Mica sensor motes. Our full-fledged nesC implementation of GPSR on the Mica motes, depicted in Figure 11, includes the GG and the RNG planarization algorithms (chosen via a configuration parameter), as well as greedy- and perimeter-mode packet forwarding³

3.3.1 Testbeds and Experiments

We deployed our GPSR implementation on a wireless sensor node testbed. This testbed (which we label R) contains 75 Mica-2 “dot” nodes deployed roughly one per room on one floor of a campus building. At the default power setting on this testbed, all nodes were within two hops of each other. To generate an interesting multi-hop topology, we reduced the CC1000 radio’s output power from 15 to 2. In the resulting topology, the average path length was around 5 hops, and the average node degree was 5.2. Note that controlling the transmit power is roughly equivalent to appropriately scaling the geographic dimensions of the testbed. Fi-

³Our implementation also includes the revised face traversal rules described in a later section. Thus, all routing failures described here can be attributed to planarization failures, and not to incorrect face traversal algorithms.

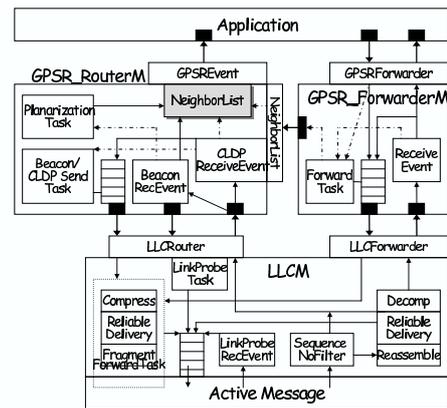


Figure 11: Implementation of GPSR.

nally, in our experiments, we statically configured nodes with their locations and started all the nodes roughly simultaneously.

We measured the fraction of all pairs of nodes on this network that could reach one another. In these measurements, we iterate over all nodes in the network, allowing one node at a time to send traffic to each other node in the network.

	topology 1 (T_1)	topology 2 (T_2)
GPSR’s GG	68.2%	97.2%
GPSR’s GG/MWP	87.8%	100%
GPSR’s CLDP	100%	100%

Table 2: Routing Success Rate in R.

3.3.2 Results

Table 2 shows results of experiments done in two different subsets of nodes on testbed R: a 50-node subset (T_1) and a 23-node subset (T_2). Notice that the underlying radio connectivity includes many crossing links (Figure 12), particularly in the dense region of the network toward the left.

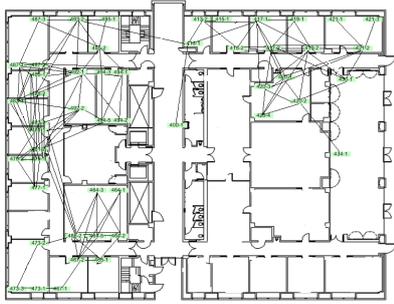


Figure 12: Node layout of $T1$.

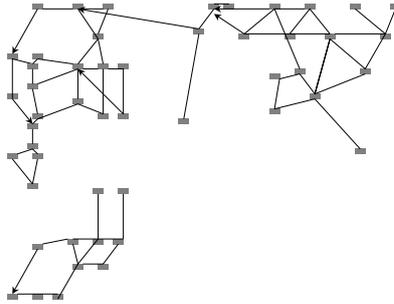


Figure 13: GPSR's GG in $T1$.

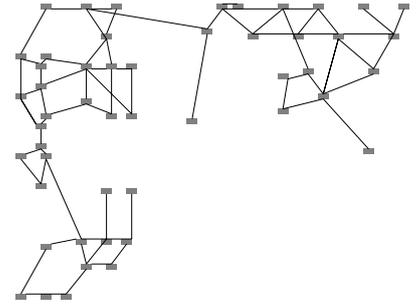


Figure 14: GPSR's GG/MWP of $T1$.

GPSR'GG: Quite surprisingly, *only 68.2% of node pairs could communicate successfully in the $T1$ testbed*. To understand why, consider the GG-derived subgraph shown in Figure 13. This subgraph clearly depicts the various kinds of planarization failures described above. The graph is clearly partitioned. In addition, unidirectional and cross links also exist in the GG-derived subgraph.

GPSR'GG/MWP: GPSR augmented with mutual witness shows slightly better performance on $T1$: 87.8% of node pairs are connected, though more than 10% of node pairs remain *persistently* disconnected. Even though disconnected links and unidirectional links are eliminated (Figure 14), a cross link is left there. Moreover, mutual witness introduces a *collinear link* in the subgraph that is described in detail in section 4.2 and Figure 20.

3.4 Cross-link Detection Protocol (CLDP)

Motivated by the failures exhibited by GPSR when run over GG- and RNG-generated planarized graphs, we have developed a protocol called CLDP, which, given an *arbitrary* connected graph, produces a subgraph on which face traversal cannot fail. Table 2 reflects the 100% success achieved by GPSR with CLDP on real networks $T1$ and $T2$.

This protocol is described in detail and proven correct elsewhere [9]; in this section, we briefly describe CLDP for completeness.

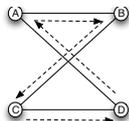


Figure 15: CLDP example.

The high-level idea behind CLDP is simple: each node, in an entirely distributed fashion, *probes* each of its links to see if it is *crossed* (in a geographic sense) by one or more other links. A probe initially contains the locations of the endpoints of the link being probed, and traverses the graph using the right-hand rule. For example, in Figure 15, consider a probe originated by node D for the link (D,A) . It contains the geographic coordinates for D and A , and traverses the graph using the right-hand rule, as shown by the dashed arrows. When the probe is about to traverse the edge (B,C) , node B “notices” that this traversal would cross (D,A) ; B records this fact in the probe so that when the probe returns to D , D notices a cross-link and “deletes” either the (A,D) link or the (B,C) link (after a message exchange with B or with C). By symmetry, the cross-links would have been detected by a probe of (A,D) originated by A or a probe of (B,C) originated either by B or C .

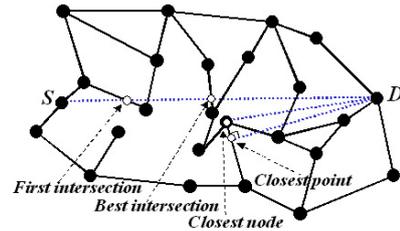


Figure 16: Points where a face change is performed.

4. FACE TRAVERSAL IN REAL ENVIRONMENTS

In this section, we consider two primitives used in face routing: face changes and the right-hand rule. We first demonstrate by example that these primitives do not work correctly in realistic wireless networks where radio ranges are non-uniform and where links may be collinear. We then describe how to make these primitives robust in realistic wireless networks, and evaluate these proposed improvements.

4.1 Incorrect Face Changes

Today's geographic routing protocols use one of four face-change rules, depicted in Figure 16:

- **Best intersection (Compass [12], AFR [14]):** After a face is entirely toured using the right-hand rule, a face change occurs at the link whose crossing of \overline{SD} is the *closest* such crossing to D on the current face.
- **First intersection (GPSR, GFG):** While walking a face using the right-hand rule, upon encountering a link that crosses the line segment \overline{SD} at a point closer to D than the point at which the current face was entered, a face change occurs at that link.
- **Closest-node other face routing (GOAFR+):** After a face is entirely toured, a face change occurs at the *node* closest to D on the current face.
- **Closest-point other face routing:**⁴ After a face is entirely toured, a face change occurs at the closest *point* to D on the current face.

Under all four of these face-change rules, when a face is toured in its entirety without discovering a point for a face change, the traversal fails, and the packet is dropped. In a planar graph, this situation

⁴In Section 4.2 of [16], Kuhn *et al.* describe this rule. But in later work, they use *closest-node* face changes in GOAFR [16] and GOAFR+ [13], as they assume network topologies that are unit-disk graphs.

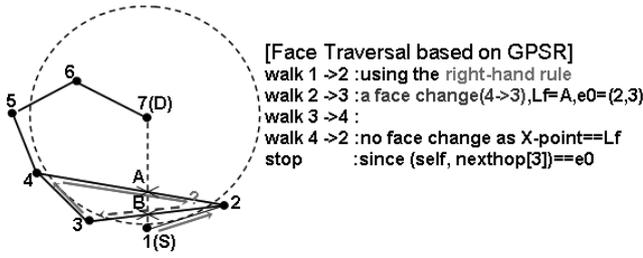


Figure 17: Failure of GPSR.

should only occur when the destination is disconnected. That is, if the destination is connected, the face-change rule must always find a face closer to D than the current face along line segment SD . We term this requirement *forward progress*.

We have found that in real wireless networks, the “first intersection” and “closest-node other face routing” rules don’t always work correctly; that is, there exist topologies on which they do not obey the forward progress requirement. We begin by illustrating these failures by example.

Figure 17 shows an example of a case (on a topology with irregular radio ranges) where GPSR’s first-intersection rule causes a packet delivery failure. Node 1 sends a packet to node 7, which is located beyond node 1’s radio range. Face traversal is invoked at node 1 as it is a local minimum to node 7. By the right-hand rule, node 2 initially chooses node 4 as its next hop. However, as link (2,4) is intersected by line (1,7) and $|A7| < |17|$, node 2 performs a face change, records location(A) in the packet’s L_f field (which records the point at which the packet first encountered a local minimum), and then forwards the packet to node 3. Thereafter, the traversal reaches node 4. Although link (4,2) intersects line (1,7), node 4 forwards the packet to node 2 as $|A7| = |L_f7|$. Next, by the right-hand rule, node 2 chooses node 3 as its next hop. At this point the traversal stops, as the choice of node 3 as the next hop completes a loop (2-3-4-2-3 ...). Thus, on this example topology, although a path to node 7 exists, GPSR’s face change rule fails to find it.

We also encountered a counterexample topology that demonstrates a failure for GFG’s face change rule, which follows a different first intersection traversal than GPSR. We elide that counterexample because of space constraints. We note further that there also exist corner-case connected unit-disk graphs on which GPSR’s and GFG’s first-intersection rules fail to find a route. We must elide these details as well because of space constraints.⁵

Figure 18 shows how GOAFR+’s “closest-node other face routing” rule can fail in a network where radio ranges are irregular. Node 1 sends a packet to node 10. Face traversal is invoked at node 1, which has no neighbor closer to node 10 than itself. Although a face (1-2-3-4-5-6-7-1) is toured in its entirety, and a closer intersection to node 10 is detected, the face traversal doesn’t encounter a *node* closer than node 1 to the destination—that is, it fails, even though paths to the destination exist.

⁵Note that it is unclear whether there exists a first-intersection face-change rule immune to these failures, either on unit-disk or non-unit-disk graphs. Such a rule is of little practical interest, however, given the relatively comparable stretch of best-intersection vs. first-intersection when combined with greedy routing in practice (see Section 4.4).

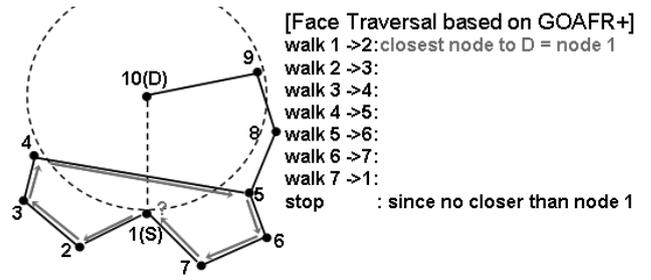


Figure 18: Failure of GOAFR+.

4.2 The Right-hand Rule and Collinear Links

Collinear links occur, as the term suggests, when portions of two or more links fall exactly on top of one another⁶. Intuitively, such links present an ambiguity for the right-hand rule:

- If the right-hand rule identifies more than one candidate next hop, which is chosen among them?
- Does the right-hand rule recognize an angle between two collinear links as 0 or 2π ?

We will show later in this section that face traversal in the presence of these ambiguities can result in delivery failures.

Under the unit-graph assumption, collinear links do not occur in a planarized network subgraph, as nodes at the endpoints of collinear links are sufficiently close to be connected to one another, and the planarization removes such *regular collinear links*. In real wireless networks, where the unit-graph assumption is routinely violated, however, there may exist *irregular collinear links* that cannot be removed by planarization. Figure 19 shows example topologies that contain regular and irregular collinear links.

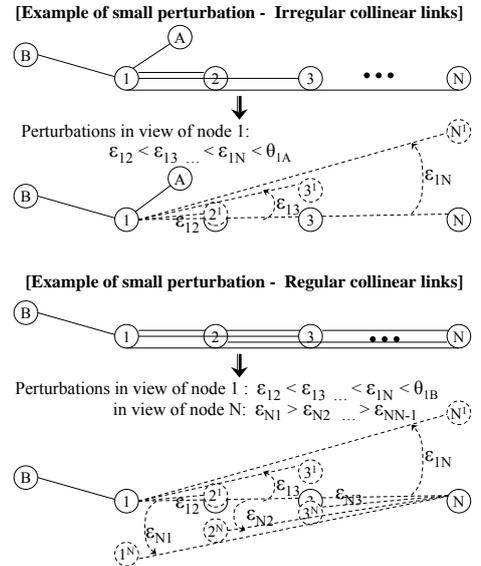


Figure 19: Examples of small perturbation.

Planarization converts irregular collinear links into unidirectional links. For example, consider the upper right of Figure 20, in which node 3 holds inaccurate information concerning its own position. At node 2, the GG planarization removes link (2,4), as

⁶Collinear links do occur in practice. They are often caused by the limited resolution of node coordinates. We observed one collinear link in our GPSR/GG/MWP experiments (Figure 14) on the Berkeley testbed.

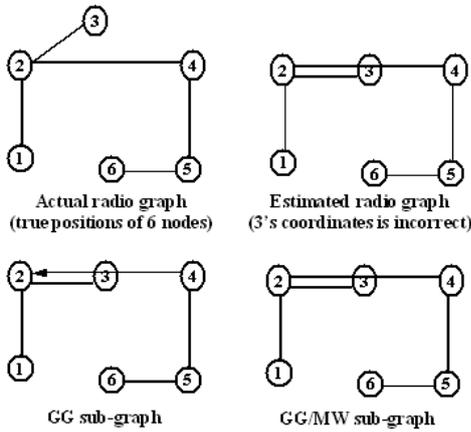


Figure 20: Pathologies by collinear links.

node 3 is a witness for the link. However, at node 4, the GG planarization keeps link (4,2), as node 4 can't see node 3. Thus, a unidirectional link occurs in the GG planarization, as shown in the lower-left of Figure 20.

As described in Section 3, planarizing a topology robustly requires using either mutual witness (Section 3.2) or CLDP (Section 3.4). However, both of these techniques can produce collinear links, as shown in the lower-right of Figure 20.

Let us now use examples to examine the ambiguity of the right-hand rule in the presence of collinear links in more detail. Suppose node 1 in the lower right of Figure 20 sends a packet to node 6. Face traversal for the packet begins at node 1, as node 1 is a local minimum with respect to node 6. When the traversal reaches node 2, which node should be the next hop, node 3 or node 4? If node 4 is chosen, the traversal can reach the destination via node 4. Suppose, however, that link (3,6) were added and link (5,6) were removed. In that case, forwarding via node 4 cannot reach the destination. Returning to the topology pictured in the lower right of Figure 20, if node 3 were chosen as the next hop, the traversal will return the packet to node 2, as node 3's only link is (3, 2). When the packet returns to node 2, another ambiguity rears its head: *is the angle from link (2,3) to link (2,4) 0 or 2π ?* If it is 2π , node 1 is chosen as the next hop, in which case the traversal stops after looping back to node 1. If the angle is 0, node 4 is chosen as the next hop and the traversal can reach the destination. However, it is not always possible to consistently assume a choice of 0 for the angle between two collinear links. Suppose that link (6,5) does not exist, but other multi-hop paths between nodes 6 and 2 exist. In this case, the traversal can't reach the destination, as the packet will loop infinitely along the path (2-4-5-4-2-3-2-4-5).

4.3 Robust Face Traversal

We now describe how to make face traversal robust in real wireless networks.

4.3.1 Robust Face Changes

As long as the destination is connected, the two following rules always discover a link where a face change will be performed, since both these rules result in a complete tour of the face, and preserve the *forward progress* requirement. Using simulation, we confirm in Section 4.4 that geographic routing using either of these face change rules never fails.

Best Intersection Face Changes : As proven in [14], best-intersection face changes must make progress toward the destination in a planar graph. That is, the best intersection rule for face

change always finds a route to a connected destination, even if the network's topology violates the unit-disk assumption, so long as the topology is planar.

Closest Point Face Changes [16] : Also, as proven in [16], the closest point rule for face change always finds a route to a connected destination.

4.3.2 Revised Right-Hand Rule

To resolve the ambiguity caused by collinear links, we now introduce the concept of *small perturbation* of node positions. Suppose a node has multiple neighbors, the links to which are collinear, and that the set of all collinear links is L_c .

(1) Qualitatively, we perturb the coordinates of each neighbor so that the link l_i to that neighbor is rotated counterclockwise from its initial (collinear) position. Each link is rotated a different amount, to eliminate collinearity of links.

(2) Quantitatively, the relative magnitudes of the rotations (ϵ_i) of the links are determined by the relative magnitudes of the Euclidean lengths ($|l_i|$) of the links. That is, the longer the link, the greater its rotation. In addition, a collinear link should be rotated an angle (ϵ_i) less than the minimum angle (θ_{\min}) between the link's original site and any links to neighbors that are counterclockwise from the link's original site. That is, $\forall i, j \in L_c, i \neq j, |l_i| > |l_j| \Leftrightarrow 0 < \epsilon_j < \epsilon_i < \theta_{\min}$.

As proven in Theorem A.2, face routing with small perturbations always works in the presence of collinear links.

Practical Collinear Link Robustness

Even though small perturbations eliminate collinearity, they are difficult to implement in practice, for reasons including the poor numerical stability of representing small ϵ_i , &c. Thus, we instead present an alternative modification to the right-hand rule that clearly answers the two questions raised in Section 4.2. Suppose that during face traversal, a node with collinear links is reached.

(1) If the traversal comes from link (l_i) among the set of collinear links L_c , the right-hand rule treats each (l_j) among the collinear links as follows: If the Euclidean length ($|l_i|$) of l_i is shorter than the Euclidean length ($|l_j|$) of l_j , then the angle between these two links is interpreted as 0; otherwise, the angle is interpreted as 2π . That is, $(\exists l_i \in L_c) \rightarrow (\forall l_j \in L_c, (i) \angle l_i l_j = 0 \text{ if } d_{l_i} < d_{l_j}, (ii) \angle l_i l_j = 2\pi \text{ if } d_{l_i} > d_{l_j})$.

(2) If some or all among the collinear links are next-hop candidates (C) for the right-hand rule, the link (l_n) with the shortest Euclidean length is chosen as the next hop. That is, $l_n = \min_{l_i \in C} |l_i|$.

As proven in Theorem A.3, this alternative is equivalent to the above small perturbation technique.

4.4 Simulation-based Evaluation of Face Routing Improvements

To evaluate the correctness and performance of the enhancements we have proposed to the face change and right-hand rule algorithms, we use the simulation environment from [9]. We evaluate four kinds of face-change rules under TOSSIM. *FR-BI* denotes face routing using face changes at the best intersection. *FR-FI* denotes face routing using face changes at the first intersection. *OFR* denotes face routing using face changes at the closest *node* to the destination. *OFR** denotes face routing using face changes at the closest *point* to the destination. Also, we compared the combination of greedy geographic routing with these four face-change rules. *GFRB* denotes greedy routing combined with *FR-BI*; note that this combination has not been evaluated in the literature previously to our knowledge. *GPSR* denotes greedy routing combined with *FR-FI*. *GOFR+* denotes greedy routing combined with *OFR*;

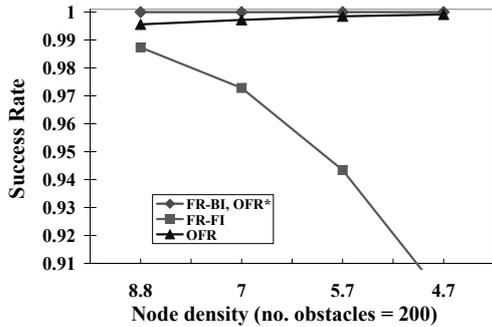


Figure 21: Success rate of four kinds of face routing.

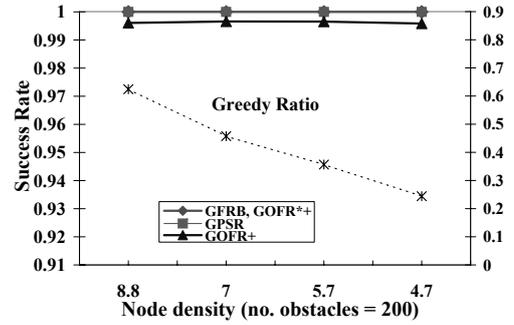


Figure 22: Success rate of four kinds of face routing combined with greedy routing.

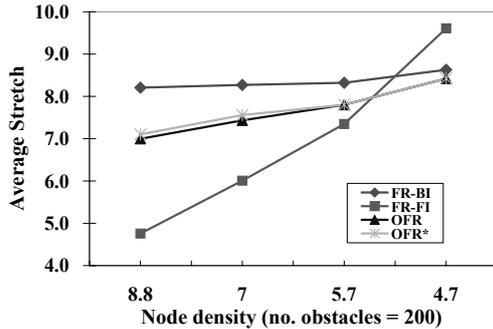


Figure 23: Stretch of four kinds of face routing.

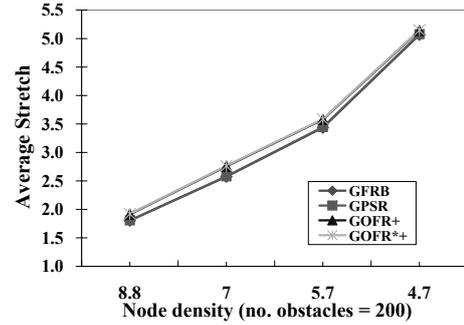


Figure 24: Stretch of four kinds of face routing combined with greedy routing.

this algorithm is essentially *GOAFR+* [13] with an infinite boundary. *GOFR*+* denotes *GOFR+* using *OFR** instead of *OFR*.

In our simulations, we consider two performance measures: *Success rate* measures the fraction of sender/receiver pairs for which packet transmissions from the sender are successfully received. *Average stretch* measures the average path stretch across all sender/receiver pairs. The stretch of a path between two nodes is the ratio of the number of hops in the path to the number of hops in the shortest path between those two nodes.

In each simulation, we place nodes uniformly at random, and then run CLDP to generate a planar network subgraph. Next, we send packets pairwise bidirectionally between nodes. Each data point in the results we present is the mean across 50 randomly generated topologies, and we have verified that this is sufficient to produce negligible 95% confidence intervals for the mean values of our metrics.

4.4.1 Success Rate

Figure 21 shows the packet delivery success rate as a function of node density for pure face routing protocols (without greedy routing), in the presence of 200 randomly placed radio-opaque obstacles.⁷ Our measure of node density is the average number of neighbors of a node.

As expected, *FR-BI* and *OFR** allow perfect delivery across all node densities we evaluated. Interestingly, the *FR-FI* algorithm offers poor success rates. At most densities, *OFR* establishes paths between 99.5% of node pairs, but not between all pairs of nodes.

Figure 22 shows success rate as a function of node density for face routing protocols combined with greedy routing. As expected, *GFRB* and *GOFR*+* always succeed, given that *FR-BI* and *OFR** alone always succeed. *GPSR* achieves 100% success, despite the low success rate of *FR-FI* alone; greedy routing’s early exit from

perimeter mode seems to mask *FR-FI*’s poor behavior. The success rate for *GOFR+* is similar to that for *OFR*—that is, greedy routing does not appear to recover from the routing failures caused by *OFR*.

4.4.2 Average Stretch

Figure 23 plots average path stretch as a function of node density for pure face routing protocols, in the presence of 200 obstacles. As expected, at most densities, *FR-FI* exhibits lower stretch than other alternatives that always tour each face in its entirety. However, at the lowest node densities, *FR-FI* exhibits the greatest stretch. Figure 24 plots the average stretch as a function of node density for face routing protocols combined with greedy routing. Surprisingly, the stretches of *GFRB* and *GOFR*+* approach that of *GPSR*, despite the longer stretches of pure *FR-BI* vs. pure *FR-FI* and of *OFR** vs. pure *FR-FI*. Note that a significant fraction of packets encounter perimeter forwarding, and cannot be delivered using greedy forwarding alone; the “greedy ratio” curve in the figure shows (on the right y-axis) the fraction of packets delivered using only greedy-mode forwarding.

Interestingly, the stretch of all compared protocols is similar because of the *fallback-to-greedy* mechanism⁸—that is, when a packet in perimeter mode visits a node closer to the destination than that at which it entered perimeter mode, it is returned to greedy mode.

In summary, geographic routing with either the best intersection rule or the closest point rule allows perfect delivery and exhibits good average stretch.

⁸There are two previously proposed ways to do so, given a local minimum at m . *GFRB* and *GPSR* fall back to greedy mode at the first node closer to the destination than m . *GOAFR+* considers falling back to greedy mode when considering a face change, and also after counters of nodes visited closer to m than the destination and closer to the destination than m reach a particular ratio.

⁷The obstacle model we use is identical to that used in [9].

5. CONCLUSION

We have cataloged the geographic routing pathologies caused by radio range irregularities and imprecise localization, in the form of planarization failures and face traversal failures. Our evaluation in deployment and simulation has confirmed that these pathologies can cause existing geographic routing protocols to fail persistently for particular node pairs. It also confirms that by combining geographic routing techniques designed specifically for correctness without requiring the unit-disk assumption, including CLDP, the revised right-hand rule, and robust face change rules, geographic routing can be made to work correctly in real wireless networks.

6. REFERENCES

- [1] BOSE, P., MORIN, P., STOJMENOVIC, I., AND URRUTIA, J. Routing with guaranteed delivery in ad hoc wireless networks. In *Proc. ACM DIALM Workshop* (Seattle, WA, USA, Aug. 1999), pp. 48–55.
- [2] FINN, G. Routing and addressing problems in large metropolitan-scale internetworks. Tech. Rep. ISI/RR-87-180, USC/Information Sciences Institute, Mar. 1987.
- [3] GABRIEL, K., AND SOKAL, R. A new statistical approach to geographic variation analysis. *Systematic Zoology* 18 (1969), 259–278.
- [4] GAO, J., GUIBAS, L., HERSHBERGER, J., ZHANG, L., AND ZHU, A. Geometric spanner for routing in mobile networks. In *Proc. ACM MobiHoc* (Oct. 2001), pp. 45–55.
- [5] HILL, J., SZEWCZYK, R., WOO, A., HOLLAR, S., CULLER, D., AND PISTER, K. System architecture directions for networked sensors. In *Proc. 9th ACM ASPLOS* (Cambridge, MA, USA, Nov. 2000).
- [6] KARP, B. *Geographic Routing for Wireless Networks*. PhD thesis, Harvard University, 2000.
- [7] KARP, B. Challenges in geographic routing: Sparse networks, obstacles, and traffic provisioning. Presentation at the DIMACS Workshop on Pervasive Networking, May 2001.
- [8] KARP, B., AND KUNG, H. T. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proc. ACM/IEEE MobiCom* (Boston, Mass., USA, Aug. 2000), ACM, pp. 243–254.
- [9] KIM, Y.-J., GOVINDAN, R., KARP, B., AND SHENKER, S. Geographic routing made practical. In *Proc. USENIX Symposium on Network Systems Design and Implementation* (Boston, Massachusetts, USA, May 2005), USENIX.
- [10] KLEINROCK, L., AND TAKAGI, H. Optimal transmission ranges for randomly distributed packet radio terminals. *IEEE Trans. Comm.* 32, 3 (1984), 246–257.
- [11] KO, Y.-B., AND VAIDYA, N. Location-aided routing in mobile ad hoc networks. In *Proc. ACM/IEEE MobiCom* (Aug. 1998).
- [12] KRANAKIS, E., SINGH, H., AND URRUTIA, J. Compass routing on geometric networks. In *Proc. 11th Canadian Conference on Computational Geometry*, Aug. 1999.
- [13] KUHN, F., WATTENHOFER, R., ZHANG, Y., AND ZOLLINGER, A. Geometric ad-hoc routing: Of theory and practice. In *Proc. ACM PODC* (Boston, MA, USA, July 2003).
- [14] KUHN, F., WATTENHOFER, R., AND ZOLLINGER, A. Asymptotically optimal geometric mobile ad-hoc routing. In *Proc. ACM DIALM POMC Workshop* (Sept. 2002).
- [15] KUHN, F., WATTENHOFER, R., AND ZOLLINGER, A. Ad-hoc networks beyond unit disk graphs. In *Proc. ACM DIALM POMC Workshop* (Sept. 2003).
- [16] KUHN, F., WATTENHOFER, R., AND ZOLLINGER, A. Worst-case optimal and average-case efficient geometric ad-hoc routing. In *Proc. ACM MobiHoc* (2003).
- [17] LI, X., KIM, Y. J., GOVINDAN, R., AND HONG, W. Multi-dimensional range queries in sensor networks. In *Proc. ACM Sensys* (Nov. 2003).
- [18] NEWSOME, J., AND SONG, D. GEM: Graph embedding for routing and data-centric storage in sensor networks with geographic information. In *Proc. ACM Sensys* (Nov. 2003).
- [19] RAO, A., RATNASAMY, S., SHENKER, S., AND STOICA, I. Geographic routing without location information. In *Proc. ACM/IEEE MobiCom* (Oct. 2003), pp. 96–108.
- [20] RATNASAMY, S., KARP, B., YIN, L., YU, F., ESTRIN, D., GOVINDAN, R., AND SHENKER, S. GHT: A geographic hash table for data-centric storage. In *Proc. ACM WSNA Workshop* (Atlanta, Georgia, USA, Sept. 2002), ACM, pp. 78–87.
- [21] SAVVIDES, A., HAN, C.-C., AND SRIVASTAVA, M. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proc. ACM/IEEE MobiCom* (Rome, Italy, July 2001), ACM, p. to appear.
- [22] SEADA, K., HELMY, A., AND GOVINDAN, R. Localization errors on geographic face routing in sensor networks. In *Proc. IEEE IPSN Workshop* (Berkeley, CA, USA, Apr. 2004).
- [23] TOUSSAINT, G. The relative neighborhood graph of a finite planar set. *Pattern Recognition* 12, 4 (1980), 261–268.
- [24] WOO, A., TONG, T., AND CULLER, D. Taming the underlying challenges of reliable multihop routing. In *Proc. ACM Sensys* (Los Angeles, CA, Nov. 2003).
- [25] ZHAO, J., AND GOVINDAN, R. Understanding packet delivery performance in dense wireless sensor networks. In *Proc. ACM Sensys* (Los Angeles, CA, Nov. 2003).
- [26] ZHOU, G., HE, T., KRISHNAMURTHY, S., AND STANKOVIC, J. A. Impact of radio irregularity on wireless sensor networks. In *Proc. ACM Mobisys* (June 2004).

APPENDIX

A. LEMMA AND THEOREMS

Lemma A.1 *CLDP will not produce regular collinear links when run on a graph to which small perturbations have been applied.*

PROOF. Consider n collinear nodes: $node_i, node_{i+1}, \dots, node_{i+n-1}$ for $n > 2$, with $node_{i+k}$ between $node_{i+k-1}$ and $node_{i+k+1}$ for $0 < k < n - 1$; $node_{i+k}$ has $n-1$ links to all other collinear nodes.

Suppose CLDP probes a link from $node_{i+k}$ to $node_{i+k+1}$. When CLDP reaches $node_{i+k+1}$, the right-hand rule selects $node_{i+k-1}$ as the next hop, since $node_{i+k+1}$ has a link to $node_{i+k-1}$ and that small perturbation is the next link counter-clockwise. CLDP on $node_{i+k+1}$ detects cross-links: $(node_{i+k}, node_{i+k+1})$ and $(node_{i+k+1}, node_{i+k-1})$. The CLDP probe finally returns to $node_{i+k}$. Thereafter, CLDP removes the link $(node_{i+k+1}, node_{i+k-1})$ from the routable subgraph. That is, CLDP removes every 2-hop-long link overlapped by link $(node_{i+k}, node_{i+k+1})$, including link $(node_{i+2}, node_i)$, link $(node_{i+3}, node_{i+1})$, \dots , link $(node_{i+n-1}, node_{i+n-3})$. Subsequently, when CLDP probes the link from $node_{i+k}$ to $node_{i+k+1}$, the protocol removes every 3-hop-long link overlapped by link $(node_{i+k}, node_{i+k+1})$, including link $(node_{i+3}, node_i)$, link $(node_{i+4}, node_{i+1})$, \dots , link $(node_{i+n-1}, node_{i+n-4})$. Through repeated probing, CLDP continues to remove overlapped links. Finally, after $n - 2$ CLDP probing repetitions, every link overlapped by link $(node_{i+k}, node_{i+k+1})$ has been eliminated from the resulting subgraph. \square

Theorem A.2 *Face routing with small perturbations routes correctly in a subgraph produced by CLDP.*

PROOF. Lemma A.1 together with Theorem 4.2 in [9] prove that CLDP does not partition a connected graph; only leaves cross-links in a graph that do not cause face routing to fail; and removes regular collinear links from a graph. Also, it has been proven in [9] that face routing works correctly on a graph that contains no irregular collinear links.

Hence, we prove that face routing works correctly on a graph containing irregular collinear links. Even though some irregular collinear links are left in a CLDP subgraph, small perturbation removes any collinear links in the subgraph, as those links are rotated so as not to be collinear. Hence, face traversal does not encounter ambiguities associated with collinear links. \square

Theorem A.3 *Face routing using practical collinear link robustness is equivalent to face routing with small perturbations.*

PROOF. First, we prove that small perturbation implies practical collinear link robustness (PCLR). Suppose face traversal reaches a node with collinear links.

(I) Suppose the traversal arrives on a collinear link ($l_i \in L_c$). The right-hand rule after small perturbation treats each collinear link ($l_j \in L_c$) as follows: l_i is rotated from its original site by an angle of ε_i and l_j is rotated by ε_j . Hence, the right-hand rule recognizes $\angle l_i l_j$ as $\varepsilon_j - \varepsilon_i$: if $|l_i| < |l_j|$, $\angle l_i l_j = \varepsilon_j - \varepsilon_i \approx 0$, since $0 < \varepsilon_i < \varepsilon_j$; if $|l_i| > |l_j|$, $\angle l_i l_j = 2\pi + (\varepsilon_j - \varepsilon_i) \approx 2\pi$, since $0 < \varepsilon_j < \varepsilon_i$. This result corresponds to the first part of the definition of PCLR.

(II) Suppose the traversal arrives on a non-collinear link l_i . The right-hand rule after small perturbation treats each collinear link ($l_j \in L_c$) as follows: l_j has been rotated from its original site l_{j_0} by an angle of ε_j . Call the angle between l_i and l_{j_0} θ . Hence, the right-hand rule recognizes $\angle l_i l_j$ as $\varepsilon_j - \theta$: $\angle l_i l_{j_0} = 2\pi + (\varepsilon_j - \theta)$, since $0 < \varepsilon_j < \theta$.

If a collinear link in L_c is one of the set of next-hop candidates C , in both cases (I) and (II), a link whose ε_j is the lowest among the candidates is chosen as the next-hop, as the right-hand rule by definition chooses a link that minimizes $\angle l_i l_j$. Suppose the chosen link is $l_n \in L_c$. The ε_n of link l_n is the lowest among next-hop candidates C . Thus, the Euclidean length ($|l_n|$) of link l_n is the least among next-hop candidates C , according to the definition of small perturbation. This result corresponds to the second part of the definition of PCLR.

Next, we prove that PCLR implies small perturbation. Suppose face traversal reaches a node with collinear links. PCLR discriminates among collinear links according to the Euclidean lengths of each link. This discrimination is qualitatively the same as small perturbation's rotation of collinear links according to their Euclidean length.

We now show this equivalence quantitatively. If a collinear link is a next-hop candidate for the right-hand rule, for a packet arriving on link l_i , the chosen next-hop is a link (l_n) whose length is shortest among all candidates. That is, $\exists l_n \in L_c$ such that $\forall l_j \in L_c, l_n \neq l_j, |l_n| < |l_j|$. This means that $\angle l_i l_n < \angle l_i l_j$.

If the traversal arrives along a collinear link ($l_i \in L_c$), then

$$\angle l_i l_n = \varepsilon_n - \varepsilon_i \quad (1)$$

and

$$\angle l_i l_j = \varepsilon_j - \varepsilon_i. \quad (2)$$

Thus, $\varepsilon_n < \varepsilon_j$, since $\angle l_i l_n < \angle l_i l_j$.

If the traversal arrives along a non-collinear link and the angle between the ingress link and the collinear links is θ , then

$$\angle l_i l_n = 2\pi + (\varepsilon_n - \theta) \quad (3)$$

and

$$\angle l_i l_j = 2\pi + (\varepsilon_j - \theta). \quad (4)$$

Thus, $(\varepsilon_j - \theta) < 0$ since $2\pi + (\varepsilon_j - \theta) < 2\pi$. Also, $\varepsilon_n < \varepsilon_j$ since $\angle l_i l_n < \angle l_i l_j$. As a result, $\varepsilon_n < \varepsilon_j < \theta$, according to equations 1, 2, 3, and 4. \square