Available online at www.sciencedirect.com

ScienceDirect

Ad Hoc Networks 5 (2007) 855–871

www.elsevier.com/locate/adhoc

Modeling and analyzing the correctness of geographic face routing under realistic conditions

Karim Seada [a,*], Ahmed Helmy [b], Ramesh Govindan [c]

[a] *Nokia Research Center, 955 Page Mill Road, Palo Alto, CA 94304, United States*
[b] *Computer and Information Science and Engineering Department, University of Florida, Gainesville, United States*
[c] *Computer Science Department, University of Southern California, Los Angeles, United States*

Received 16 February 2007; received in revised form 17 February 2007; accepted 20 February 2007
Available online 1 March 2007

Abstract

Geographic protocols are very promising for wireless ad hoc and sensor networks due to the low state storage and low message overhead. Under certain idealized conditions, geographic routing – using a combination of greedy forwarding and face routing – has been shown to work correctly and efficiently. In this work we model and analyze the correctness of geographic routing under non-ideal realistic conditions. We present a systematic methodology for micro-level behavioral analysis that shows that conditions that violate the unit-graph assumption of network connectivity, such as location errors, obstacles and radio irregularity, cause failure in planarization and consequently face routing. We then discuss the limitations of fixing these failures and prove that local algorithms that use only information up to a limited number of hops are not sufficient to guarantee face routing delivery under arbitrary connectivity. In addition, we analyze the effect of location errors in more detail to identify the possible protocol error scenarios and their conditions. We present results from an extensive simulation study about the effects of location errors on GPSR and GHT to quantify their performance degradation at different error ranges, distributions and error models. Based on our analysis we present a potential fix based on local information sharing that improves the performance significantly but does not remove all failures. Finally, we conclude that in order to avoid all failures under arbitrary connectivity, we need a non-local algorithm that can search or propagate information for an unlimited number of hops.
© 2007 Elsevier B.V. All rights reserved.

Keywords: Geographic routing; Face routing; Planarization; Location errors; Wireless sensor networks; Data-centric storage

1. Introduction

Geographic routing protocols [16,20] are very attractive choices for routing in sensor networks for several reasons. First, such protocols can incur low route discovery overhead relative to flooding-based approaches, and hence conserve energy. Second, these protocols are stateless in the sense that

* Corresponding author. Tel.: +1 650 521 3128.
E-mail addresses: kseada@alumni.usc.edu (K. Seada), helmy@cise.ufl.edu (A. Helmy), ramesh@usc.edu (R. Govindan).
URLs: http://research.nokia.com/people/karim_seada (K. Seada), http://www.cise.ufl.edu/~helmy (A. Helmy), http://pollux.usc.edu/~ramesh (R. Govindan).

1570-8705/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.adhoc.2007.02.008

nodes need not maintain per-destination information, and only neighbor location information is needed to route packets. Third, they are more responsive to dynamics, since changes need to propagate for a single hop only. For these reasons, geographic routing is becoming the protocol of choice for many emerging applications in sensor networks, such as data-centric storage [23] and distributed indexing [10]. Hence, it is quite crucial to develop a detailed understanding of the behavior of geographic routing for various practical settings and to evaluate its performance and (more importantly) its correctness in those settings.

Most geographic routing protocols use greedy forwarding as its basic mode of operation, where the next forwarding hop is chosen to minimize the distance to the destination. Greedy forwarding, however, fails in the presence of *voids* or *dead-ends*. In order to provide *correct* routing in the presence of dead-ends, *face routing* has been proposed to route around the void. The most commonly used geographic routing protocols include greedy forwarding coupled with face routing.

The evaluations of geographic routing protocols have commonly assumed an ideal network connectivity graph based on the unit-graph assumption (a pair of nodes is connected if and only if the distance between them is below a certain threshold). The unit-graph assumption is valid under some ideal conditions such as the availability of accurate location information, the nonexistence of obstacles, and an ideal spherical wireless radio range. In reality these conditions are violated: obstacles do exist, experimental studies have shown that wireless channels have irregular shape and location measurements (in systems that either rely entirely on GPS, or infer location using ad-hoc localization systems [13]) are often noisy and incur some error.

In this paper, we analyze systematically the pathologies that can arise in face-routing based geographic protocols. Our methodology for this analysis is novel: using an elaborate, micro-level analysis of face routing, we infer the conditions in which the protocol correctness is violated. Our analysis shows that failures in face routing happen for two reasons: disconnections in the planar graph or cross-links. We show that these failures happen due to planarization failures when the unit-graph assumption is violated and present the conditions for that, based on the inconsistency in the distance between two nodes and the existence of a wireless link in between. This inconsistency could happen for various reasons such as location errors, obstacles and irregular radio coverage.

We then analyze one of those reasons, location errors, in more detail to identify the possible protocol error scenarios and the conditions for these scenarios to happen. We perform also extensive simulations to evaluate and quantify the effects of location errors on two protocols that use face routing; GPSR [16] and GHT [23]. Our study shows that realistic location errors can in fact lead to incorrect (non-recoverable) behavior and noticeable degradation of performance, more so for GHT than for GPSR. In some cases, more than 10% storage failure of sensor network events can occur in the presence of 10% location error. Based on our analysis and error classification, we introduce a simple protocol fix based on local information sharing that eliminates the most likely protocol errors. Though our simulations of this fix show a near perfect performance in the existence of location errors, this fix does not solve all problems.

We prove that it is not possible for a local algorithm that uses only information from neighbors, up to a limited number of hops, to guarantee correct face routing under all conditions that violate the unit graph assumption. We conclude that a non-local algorithm is necessary to guarantee delivery.

The rest of the paper is structured as follows. In Section 2 we present the related work. In Section 3 we provide the model and assumptions of our study. Section 4 explains the methodology for micro-level analysis, the errors that could happen, and the different conditions for errors. Section 5 discusses the possibility and limitations for fixing these errors. Section 6 studies the effect of location errors in more detail and presents a local fix. Section 7 contains the simulation results. Finally, the conclusions are presented in Section 8.

2. Related work and background

Early work in geographic routing considered only greedy forwarding [8] by using the locations of nodes to move the packet closer to the destination at each hop. Greedy forwarding fails when reaching a local maximum, a node that has no neighbor closer to the destination. CompassII [19] presents a face routing algorithm that guarantees message delivery on a geometric graph by traversing the edges of planar faces intersecting the line between the source and the destination. Bose et al. [4] discuss algorithms for extracting planar graphs from unit graphs and for face routing in the planar graphs to guarantee delivery. Due to the inefficient paths resulting from face routing, they propose combining face routing with greedy forwarding to improve the path length. In addition to face routing, there were also other approaches proposed in the literature to deal with the dead end problem such as [2,14]. For an extensive survey about these approaches and others refer to [25].

GPSR [16] is a geographic routing protocol for wireless networks that combines greedy forwarding and face routing (perimeter routing). Each packet contains the position of the destination and nodes need only local information about their position and their immediate neighbors' positions to forward the packet. Each node forwards the packet to the neighbor closest to the destination using greedy forwarding. When greedy forwarding fails, face routing is used to route around dead-ends until closer nodes to the destination are found.

Geographic routing has been found effective to support data-centric storage in sensor networks. GHT [23] is a geographic hash table system that hashes keys into geographic locations, and stores the key-value pair at the sensor node closest to the hash of its key. GHT uses face routing in a novel way to identify a packet home node (the node closest to the geographic destination). Packets enter face routing at the home node (since no neighbor could be closer to destination), and traverse the perimeter that enclose the destination (home perimeter) before returning back to the home node. Another system that builds on top of GHT is DIFS [10]. DIFS provides a distributed index for efficient index construction and range searches in sensor networks.

Many localization systems have been proposed in the literature: GPS, infrastructure-based localization systems [31,22], and ad-hoc localization systems [5,24,21]. We do not discuss these systems in detail; the interested reader is referred to an extensive survey of localization by Hightower and Borriello [13]. We will, however, observe that in all these localization systems an estimation error is incurred that depends on the system and the environment in which it is used. Based on our reading of the literature, we believe that a localization error of 1–10% of the radio range is very reasonable to assume even for the best existing schemes. Clearly, some systems can be more accurate (*e.g.*, GPS or dGPS based systems), but it would be prudent to ensure that face routing systems are robust to location errors that are at the higher end of this range.

In [12], simulation results were shown for the effect of localization errors on the performance of greedy forwarding. Their conclusion is that routing performance is not significantly affected when the error is less than 40% of the radio range. Face routing is not considered in that work. Our earlier work [26] was the first to point to the problems that could happen to face routing under non-ideal conditions and to show specific scenarios for face routing failure. A lot of interest has spurred in geographic routing since then with several studies following on related issues [17,18,11,29,28]. Other studies focused on more specific issues such as irregular ranges [1] and unidirectional links [7]. In the current paper we are presenting a framework for the correctness of geographic face routing, analyzing the general conditions under which failures can happen (Section 4) and considering additional issues such as obstacles and irregular radio ranges. We are also presenting the limitations for solving the errors and proofs that local algorithms are not sufficient (Section 5). Additional simulation results are provided for non-uniform distributions and correlated error models.

3. Model and assumptions

Before we analyze the impact of non-ideal conditions on geographic routing protocols, we discuss our model of the wireless network. The network consists of a set of wireless nodes, where each node knows its position using some localization technique (the precise technique used is immaterial for our purposes). All nodes have the same radio range and they broadcast beacons to their neighbors, so that each node knows about its neighbors and their locations. In an ideal environment: (i) nodes detect and announce their accurate locations, (ii) radio ranges of all nodes are exact and symmetric, (iii) there are no obstacles and so nodes within radio range can always communicate, and (iv) changes in the topology are slow comparable to the announcements such that all nodes have a consistent view of the network. Admittedly, these are idealized assumptions that do not hold in practice. As we illustrate in this paper, violation of any of these assumptions can result in routing pathologies even if the rest of these idealized assumptions hold.

The geographic routing protocol consists in general of greedy forwarding combined with face routing to overcome dead-ends. In order to perform face routing, a planar connectivity graph for the network needs to be constructed. A local planarization algorithm such as *RNG* or *GG* is used to create a planar graph for face routing. There is a class of protocols that follow this model [4,16,20].

Our study focuses on the effect of conditions that cause persistent failures on geographic routing. Thus, we assume a static and stable network (no mobility and no node failures), where nodes have consistent location information about other nodes, meaning that a node estimates its location and announces it, and all nodes observe the same estimated location for that node. In [27] we have studied the effect of probabilistic wireless channel losses on the performance of geographic routing. In the current work we are focusing on the correctness of geographic routing and we will assume that a reliable link exists to a neighbor or no link exists.

4. Micro-level analysis

In this section we present a novel methodology that detects the conditions for protocol failure such that scenarios could be generated for these failures. This methodology is based on micro-level analysis of the protocol components and using the protocol specification to detect the cases where the protocol could go wrong and the various conditions that lead to that.

Our approach starts by decomposing geographic routing into its major components, and then identifying the errors that can happen in each component and which of these errors (or combinations thereof) cause overall protocol failure. A *complete* geographic routing protocol consists of the following main components: (a) greedy forwarding, (b) planarization, and (c) face routing (also called perimeter routing or planar graph traversal). Greedy forwarding alone *does not* guarantee the delivery of packets because of dead-ends (variously called local maxima or voids). Face routing on a planar graph theoretically *does* guarantee the delivery of packets. For improved performance, face routing is typically integrated with greedy forwarding and is used as a way to overcome dead-ends when greedy forwarding fails. Wireless network connectivity is in general non-planar, this is why the planarization component is required to create a planar graph by using only a subset of the physical links during face routing.

Non-ideal conditions may cause failures in each of these components. These conditions can cause greedy forwarding to fail in forwarding a packet to a node closer to the destination. Since failures in greedy forwarding are recovered from by face routing, we shall focus our study on *persistent protocol failures* caused by failures in face routing. As we will show, face routing failures are strongly associated with planarization failures. First, we provide a more detailed view of face routing and planarization. Then we analyze face routing in more detail to detect the conditions that cause failure.

4.1. Face routing and planarization

Correct operation of face routing requires the graph to be planar. *RNG* [30] (Fig. 1) and *GG* [9] (Fig. 2) are examples of algorithms that create a planar graph from the non-planar physical topology by selecting a subset of the links and using only those links during face routing. A desirable feature in these algorithms is that they are local (a node needs to know only its own and direct neighbors' locations) and run in distributed manner, so that each node can decide the links to include for planar routing using only local information independent of other nodes. The main idea of both algorithms is for a node to exclude an edge to a neighbor from the planar graph if there is another path through a different neighbor called *witness*. The witness should exist in a specific intersection area between the two nodes of the edge. These algorithms assume a unit graph: a pair of nodes is directly connected, if and only if, the distance between them is below a certain threshold.

In face routing a packet keeps traversing planar faces getting closer to its destination. In Fig. 3,

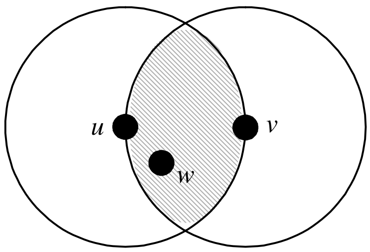


```
For each node u, where N is a list
of the neighbors of u:
 for all v ∈ N
  for all w ∈ N
   if w == v then continue
   else if d(u,v)>max[d(u,w),d(w,v)]
        remove edge (u,v)
```

Fig. 1. *RNG* planarization algorithm.

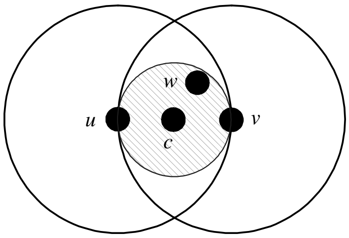


```
For each node u, where N is a list
of the neighbors of u:
 for all v ∈ N
  for all w ∈ N
   if w == v then continue
   else if d(c,w)<d(c,u) {where c
   is the midpoint of edge (u,v)}
       remove edge (u,v)
```

Fig. 2. GG planarization algorithm.

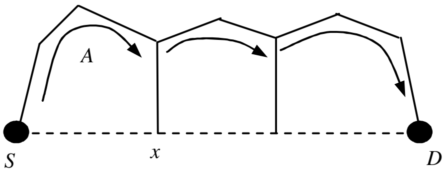

Fig. 3. Face routing.

assume node S is forwarding a packet using face routing to node D. Using the right-hand-rule the packet starts traversing face A, switching to other faces intersecting SD until reaching the face containing D. This basic mechanism is shared by all protocols that employ face routing. There are some minor differences between protocols in where to switch faces: for example, in GPSR when a packet reaches a node having an edge intersecting SD at a point closer to the destination (point x) it uses the right hand rule to enter the next face.

4.2. Error analysis

In this section we analyze the protocol components to detect the conditions for protocol failure under the current model and assumptions.

4.2.1. Face routing

Suppose node S is forwarding a packet using face routing to node D. Each face (exterior or interior) intersecting SD will have at least two intersection points. In Fig. 3, face A intersects SD at S and x. In a planar connected graph starting from S and using the right hand rule, the packet will have to reach x. The error that could happen is that the packet gets out of the face before reaching the opposite end. This can only happen due to:

(i) a disconnection in the planar graph,
(ii) or a cross-link.

Without these two errors, the packet will keep moving to closer faces until it reaches the destination face. If routing succeeds in reaching the face containing the destination point, but the destination cannot be accessed from there, then the possible reasons for that are:

(i) the destination is not at the estimated location,
(ii) nodes at the estimated location face have no access to the destination.

Since we are concerned with the correctness of the routing itself, we will focus on the first two errors (disconnections and cross-links) that cause the routing failure. We will still show a scenario for the last two errors in Section 6. Routing failure could happen due to disconnection in the planar graph or due to cross-links. By definition, a correct planarization algorithm should not disconnect the graph and should remove all cross-links, accordingly these face routing failures could only happen due to planarization failures.

4.2.2. Planarization algorithm

By looking at the algorithms in Figs. 1 and 2, the only decision the local planarization algorithm takes is whether to remove an edge from the graph or not. So the only errors during planarization are (a) to remove an edge that should not be removed, or (b) to keep an edge that should be removed.

Planarization is based only on one assumption: the unit graph assumption. If the network follows the unit-graph assumption, then planarization should succeed. Thus, for these errors to happen the unit graph assumption has to be violated.

4.2.3. Unit graph assumption

A main assumption for the correctness of planarization is that the topology is a unit graph, which means that an edge exists between two nodes if and only if the Euclidean distance between them is less than the radio range. Based on the previous analysis we can conclude that conditions that violate the unit graph assumption are the reason for planarization failure and accordingly face routing failure. In order to derive these conditions, we notice that the unit graph assumption is based on two factors: (i) the distance between two nodes and (ii) the existence of a radio link in between. The inconsistency between the distance and the radio link connectivity, in a way that violates the unit-graph-assumption, is the reason for planarization failure. There could be different reasons in the environments where the systems are deployed for this inconsistency to happen. These reasons cause the inconsistency in similar ways and accordingly lead to similar errors. Next, we will show the conditions that lead to the violation of the unit-graph assumption and examples from real-world environments on their causes.

The inconsistency between distance and connectivity can be classified into two conditions:

(i) a radio link does not exist between two nodes, even their *observed* distance is below the nominal radio range, and
(ii) a radio link does exist between two nodes, even their *observed* distance is above the nominal radio range.

In the first condition (i), when a link does not exist between two *close* nodes, this may cause disconnections in the planar graph. Some examples are shown in Figs. 4–6, where node u removes the edge uv from the planar graph while there is no other path from u to v. Since the planarization equation depends on the location of nodes, *location inaccuracy* is a reason for the error shown in Fig. 4. Node u sees node w as a witness, although its actual location is outside the intersection area and it is not connected to node v. Depending on the network planar graph, this may cause disconnection if there is no other path from node u to node v. In Fig. 5, the distance between node w and node v is below the nominal radio range, but they are not able to

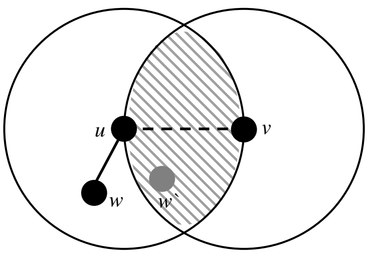

Fig. 4. Node w estimated location is in the intersection area, accordingly node u sees it as a witness and removes the edge uv. This may cause disconnection, since w actual location is outside the intersection area and it is not connected to v.

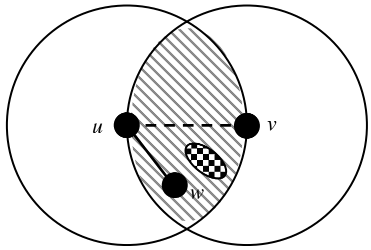

Fig. 5. The removal of edge uv by node u may cause disconnection, because there is an obstacle between nodes w and v and so they are not connected.

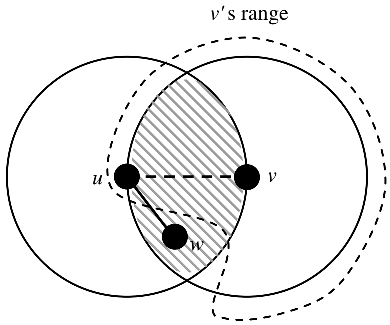

Fig. 6. The removal of edge uv by node u may cause disconnection, because node w is not in the actual radio coverage of node v.

communicate because an obstacle exists in between. *Obstacles* are a common environmental condition for violating the unit graph assumption. Node u removes the edge uv from the planar graph assuming that nodes w and v are connected. This edge removal may cause disconnection if there is no other path from node u to v. Similarly, in Fig. 6, node u removes the edge uv from the planar graph assuming that nodes w and v are connected according to a nominal radio range, but in this case they are not connected because the radio range of node v is not covering node w and so node w does not hear node v or consider it as a neighbor. *Irregular radio range* is an important condition that causes violations in the unit graph assumption which depends on the radio system and other factors in the environment. Fig. 7 is an example of how the non-existence of a link between two *close* nodes also can cause a cross-link. Because of the *obstacle*, node u does not see node w and so, if there are no other witnesses, it adds the edge uv to the planar graph, which causes a cross-link with edge xw.

The second condition (ii), a radio link does exist between two *far* nodes, can also cause cross-links in the planar graph. In Fig. 8, due to the *location*

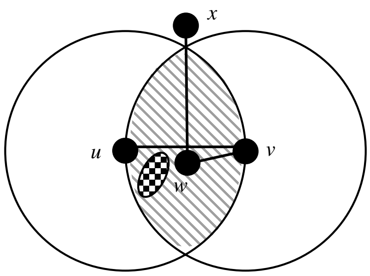

Fig. 7. Node u does not see the witness node w, because there is an obstacle in between, and so it adds edge uv to the planar graph (assuming no other witnesses exist) causing a cross-link with edge xw.

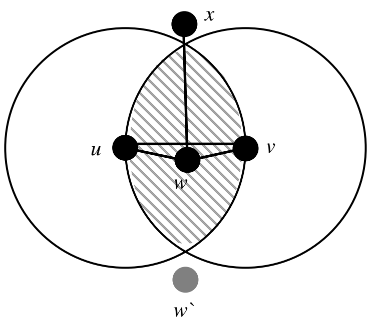

Fig. 8. Node w actual location is in the intersection area, but its estimated location is outside, accordingly node u does not consider it as a witness and it adds the edge uv to the planar graph (assuming no other witnesses exist). This causes cross-links between the edges uv and xw.

inaccuracy of node w, the observed distance between w' and x is above the nominal radio range. This causes the edges uv and xw to cross. Node u does not see w as a witness and does not remove the edge uv from the planar graph. A general case of this inconsistency is shown in Fig. 9, where the cross-links happen due to irregular radio range coverage.

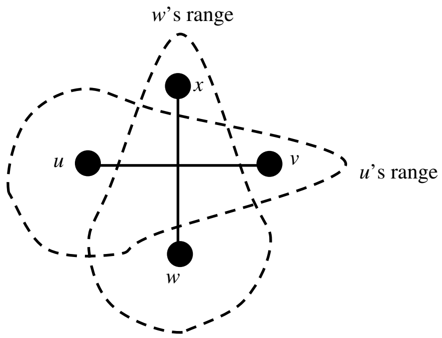

Fig. 9. Irregular radio range coverage can cause cross-links.

This could not happen under the perfect circular radio range assumption, but in reality the wireless radio coverage could take any shape, which makes this kind of cross-links a possibility.

In the previous analysis, we have shown how the two conditions of inconsistency can have different causes from real-world environments and can lead to planarization and face routing failure. The observed distance between two nodes depends on their estimated locations and accordingly is affected by the localization system. The radio link depends on the radio coverage and environmental factors that affect this coverage. Inaccurate location estimation is due to errors or inaccuracy of the localization system, and the obstacles and irregular radio range are due to the radio coverage and environmental factors that affect it. Violation of the unit graph assumption has caused one of the following errors in planarization: (i) an edge removed causing disconnection or (ii) an edge added causing cross-links. These causes are very common in many environments, but they are not the only reasons for failures. Any other factors that cause the presented two conditions of inconsistency will lead to a similar type of failures.

5. Impossibility of correct face routing using local algorithms

As we have shown by our micro-level analysis, failures in face routing can happen for two reasons: disconnections in the planar graph or cross-links. In this section, we study the possibility of having fixes that prevent face routing failures and guarantee delivery. The first question we want to answer is whether it is possible to have any algorithm (local or non-local) that fixes all the problems and guarantees delivery. If this is possible, then the second question is whether it is possible to have a local algorithm that guarantees delivery.

We will first check whether it is possible to have a planarization algorithm that can obtain a planar and connected graph from an arbitrary graph. Since face routing succeeds in planar connected graphs, the existence of such a planarization algorithm will guarantee delivery. Nonetheless, we can prove that *for arbitrary connectivity graphs, it is not possible to avoid disconnections and cross-links at the same time.* For example, in Fig. 10 adding AB and CD causes cross-links, but removing any of them causes disconnection either to node B or node D. Under

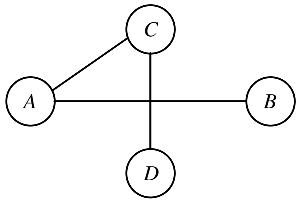

Fig. 10. Adding the edges AB and CD causes cross-links, while removing any of them causes disconnection.

the unit-graph assumption this scenario is not possible.

Proof. Suppose we have 3 separate graphs G_1, G_2, and G_3, where each graph is connected and planar. Suppose an edge e_i is added that connects G_1 and G_2, and a second edge e_j is added that connects G_1 and G_3, such that the two edges e_i and e_j are cross-links. The resulting graph G, which contains G_1, G_2, G_3, e_i, and e_j is connected. Assume that there exists a planarization algorithm (local or non-local) that can planarize G to make it connected and planar at the same time. In order for the graph G to be planar, the planarization algorithm has to remove either edge e_i or edge e_j. Since removing e_i will cause G_2 to disconnect and removing e_j will cause G_3 to disconnect, then the planarization algorithm cannot make G connected and planar at the same time, hence a contradiction. Accordingly, with arbitrary connectivity, a connected planar graph is not guaranteed using any '*planarization*' algorithm. □

However, note that even if a graph has cross-links, there is no reason to believe that face-routing on this graph will necessary fail. In some graphs it fails, but in other graphs it does not. This is different than disconnections; if a graph is disconnected then face routing will definitely fail, since two nodes will not be able to communicate. Accordingly, there is a possibility that there are algorithms that generate non-planar graphs with certain characteristics, such that face routing still can work correctly in these graphs. Recent work has demonstrated the existence of one such algorithm CLDP [17]. CLDP guarantees delivery but it needs to perform non-local searches. Each node needs to check the surrounding faces of its entire links in order to detect cross-links and decide which links to include in the face-routing graph.

In general, we would prefer to use a local algorithm. The reason is that we would like to preserve the locality characteristic of geographic routing, since it is the base for its scalability and efficiency. To clarify that, we will look at locality in more detail and observe that we will consider an algorithm local if it satisfies the following condition (we include a special condition of a single hop and a more general condition of a limited number of hops):

Locality in message propagation, which means that a node needs to know only about its direct neighbors (or at least nodes up to a limited number of hops) and does not need to search or propagate information beyond a single hop (or a limited number of hops). This is important to limit the message overhead and accordingly the power consumed. It is also extremely important in dynamic networks, so that changes in topology do not cause chains of message updates and longer convergence time. An additional property is locality in storage, which means that a node needs to store information about its direct neighbors only (or about nodes up to a limited number of hops). This is important in resource-constrained networks such as sensor networks, where the memory is limited. The storage overhead in this case depends only on the density and not on network size.

The planarization algorithms, GG and RNG, and the partial fix we will introduce in the next section are local, because all decisions for adding and removing edges in the planar graph are based only on local information.

Now we will show that a local message propagation algorithm cannot detect all cross-links in an arbitrary connectivity graph and cannot guarantee successful face routing. This can be proven by a counter example such as the scenario in Fig. 11(a).

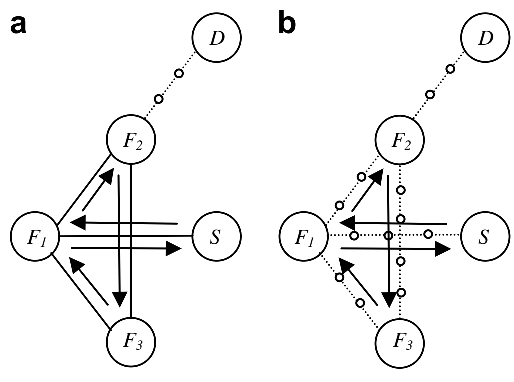

Fig. 11. Face routing to deliver a packet between S and D fails in these graphs. Dotted lines between two nodes indicates that this could be a path of multiple nodes.

In order for S to send a packet to node D, the packet will traverse S, F_1, F_2, F_3, F_1, and then returns back to S resulting in face routing failure. Node S cannot detect the cross-link by using only the information of its direct neighbors without querying multiple hops down the graph and the same thing for nodes F_1, F_2, and F_3.

In geographic routing, a local algorithm usually used to mean a single hop algorithm, but in general an algorithm that searches for a fixed number of hops in the neighborhood could also be considered local. Accordingly, we are going to generalize this result to any algorithm with a fixed number of hops.[1] By inference we can show that any search of a fixed number of hops is not sufficient to guarantee the detection of all cross-links with arbitrary connectivity. This can be shown by the example in Fig. 11(b), if we notice that the paths SF_1, F_1F_2, F_1F_3, and F_2F_3 could consist of any number of hops, and the same failure could still happen.

Proof. Suppose that two edges e_0 and e_i are cross-links and suppose that the shortest path between the two edges contain the edges $e_1, e_2, \ldots, e_{i-2}, e_{i-1}$. In order for a node on e_0 to detect the cross-link, it needs information i hops away, and since in an arbitrary graph i can grow to any value based on the network size, a limited search is not sufficient. Accordingly, we conclude that a local algorithm that searches only for a fixed number of hops cannot guarantee delivery in graphs with arbitrary connectivity. □

Based on this discussion, we summarize that *for graphs with arbitrary connectivity that violate the unit graph assumption, local algorithms that use only single-hop information (or even a fixed number of hops) cannot detect all cross-links in the graph and cannot guarantee delivery for face routing.*

In the next section, we will present a local fix that solves the most common problem that arises with location errors, which is graph disconnections, and improves the delivery significantly. This local fix will also show that the disconnection problem separately could be solved completely using a local algorithm.

[1] We mean by fixed number of hops that it is independent of the number of nodes in the network, so that the same number could be used in networks of any size. By this definition, we do not consider traversing all nodes in the network (or a factor thereof) as fixed, since it needs to increase as the network grows.

6. Detailed analysis of location errors

In Section 4 we have shown several conditions for violating the unit graph assumption. In this section, we study one of those conditions, the effect of location errors, in more detail. We present scenarios that cause protocol errors, following a systematic approach in creating the scenarios and analyzing them. This makes it easier for us to realize a complete listing of all the possible failures under the current model and assumptions. We show scenarios where only a single node has inaccurate estimated location and all other nodes are accurate. These scenarios are helpful in understanding the causes and conditions for errors under minimal discrepancy, where everything is ideal except for a single node inaccuracy. (In the next section, we study using simulation the effects of errors in random topologies, where all nodes have a random inaccurate estimated location). We then present a local fix that solves the most common problem but does not guarantee delivery. In [26] we provided quantitative analysis for the error bounds and the range of localization inaccuracy under which these errors occur.

6.1. Error scenarios

We present and discuss four main error scenarios that illustrate how errors happen for four different reasons. In these scenarios we show the accurate locations of nodes and their estimated locations. We assume that *only a single node has an inaccurate estimated location*. Even with this relatively benign assumption, routing pathologies can occur in geographic routing protocols. In our pictorial depictions of these pathologies, an edge between two nodes means that they are in range and there is a physical connection between them. In the estimated topology a dashed edge means a physical connection not included in the planar graph. The dashed circle is the accurate location of the inaccurate node. In all scenarios node S wants to forward a packet to node D and node E is the node with the inaccurate estimated location.

Fig. 12 shows a scenario where an inaccurate node location causes planarization to remove an edge that should not be removed. Node S is the closest node, among its neighbors, to node D, hence it cannot use greedy forwarding. In the accurate topology, Fig. 12(a), S uses face routing (perimeter forwarding) to forward the packet to F_1 and the packet goes around the perimeter till it reaches D.

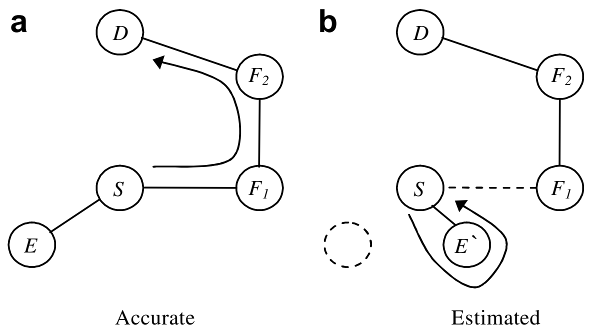

Fig. 12. Disconnection due to incorrect edge removal by planarization.

In the estimated topology, Fig. 12(b), node E has an inaccurate location such that S's planarization algorithm sees E as a witness and removes the edge (S,F_1) from the planar graph. Removal of (S,F_1) causes the planar graph to be disconnected and accordingly face routing fails to deliver the packet.

In the scenario shown in Fig. 13, the opposite happens. Edge (F_1,F_2) is removed by planarization in the accurate topology, Fig. 13(a), because node E is a witness, allowing face routing to deliver the packet from S to D. In the estimated topology, in Fig. 13(b), node E is not a witness anymore and so edge (F_1,F_2) is not removed causing edge (F_3,E) to cross it (notice that edge (F_3,E) is unidirectional). The packet loops around nodes E–F_4–F_2–F_1 until its TTL gets exhausted and it is discarded (if some kind of loop detection is used, the packet will be discarded immediately).

In Fig. 14, planarization includes all edges in both accurate and estimated topologies, but the estimated location of node E causes a cross link between (F_1,F_2) and (F_3,E) that cannot be detected

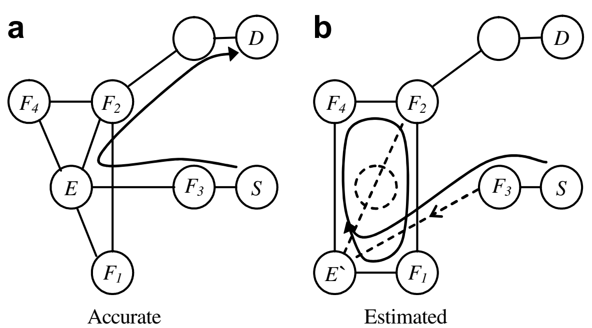

Fig. 13. Permanent loop due to planarization failure in not removing edge (notice that the figure may not be accurately scaled; distances F_3–F_2 and F_3–F_1 are longer than distances F_3–E and F_2–F_1); (a) accurate and (b) estimated.

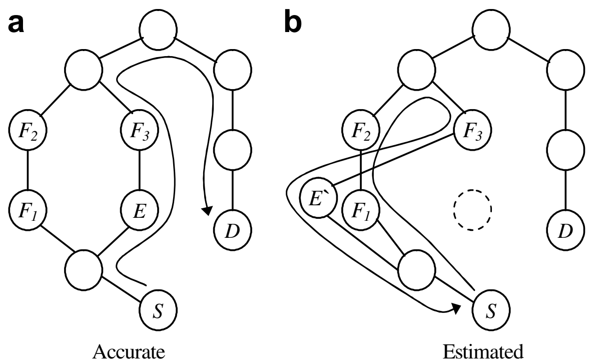

Fig. 14. Cross links causing face routing failure; (a) accurate and (b) estimated.

by the local planarization algorithm. Since face routing assumes and requires planar graphs, cross-links cause route failure.

Fig. 15 is different than the previous scenarios in that the destination (not an intermediate node) has the inaccurate location. S forwards the packet to the estimated location of D, and routing eventually succeeds in reaching the perimeter surrounding the estimated location, but since none of the nodes in that face is in range with D, the packet cannot be delivered to D. The exact sequence of nodes traversed before the packet is dropped depends on which node on the perimeter is closer to D'. In this scenario the routing itself does not fail since it reaches the announced location, but the destination is not there. A similar scenario can happen also when the destination is accurate, but other nodes not in its range have inaccurate estimated locations around its accurate position.

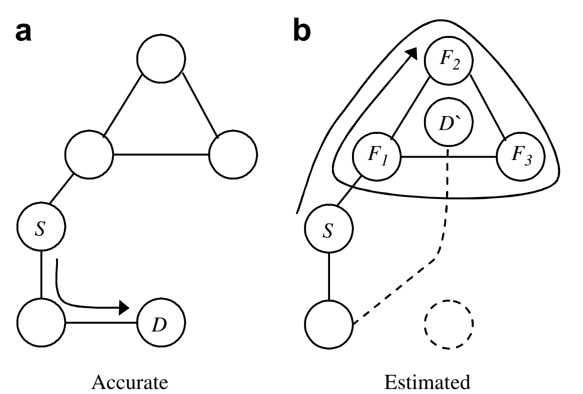

Fig. 15. Destination inaccuracy causing failure in reaching the destination. The route in this figure assumes that $F2$ is the closest node to D'. The packet is forwarded greedily to $F2$, and then face routing is used around the perimeter; (a) accurate and (b) estimated.

6.2. Partial fix

Based on our analysis, we now propose a fix that address the most likely problems and explain our rationale for choosing this fix. Not all the problems are equally likely to occur. We can categorize the problems that cause face routing failure into three main categories [26]: 1. edge removal causing disconnection, 2. cross links, and 3. inaccuracy in the destination's location. Without global knowledge about the network it is not possible to solve all the problems completely. Thus, we need to assess which problems are the most common under reasonable localization inaccuracy.

Density is a factor that affects when problems happen. At high density, greedy forwarding is used most of the time and with reasonable inaccuracy range this is unlikely to change (our simulations also show that). Since the errors happen due to face routing, dense networks look robust to these errors. So we focus on the probability of these problems in sparse networks. Intuitively, in a sparse network, disconnection seems the most serious problem that can happen. More specifically, the problem of edge removal will happen when any node enters the planarization intersection between two nodes causing their edge to be removed. This is a reasonably possible error in sparse networks even with low inaccuracy. The second problem of cross-links is very unlikely under the considered (10%) inaccuracy range as our analysis in [26] shows. Accordingly, the cross-links problem seems much less probable than the edge removal problem. In the third problem of destination inaccuracy, the destination needs to move to a face in which it is not connected to any of its neighbors, which also seems unlikely at low ranges of location inaccuracy.

Based on this analysis (and this informal analysis is supported later by our simulation results), the disconnection problem caused by edge removal seems to have the highest probability, and solving this seems likely to give the most gains in performance.[2] From the planarization algorithm, an edge is removed from the planar graph when a witness is seen by a node (e.g., in Fig. 1, node u removes edge (u,v) since there is a witness w). Disconnection happens when this witness is connected to the node removing the edge but not to the other node of the edge (w is connected to u, but not to v). Our solution for this problem is to *allow a node to remove an edge only if the other node of the edge sees the same witness* (i.e., both u and v need to see w in order for (u,v) to be removed).[3]

This fix requires modification to the planarization algorithm as follows. Before removing an edge (u,v) based on a witness w, node u needs to determine whether w is also a neighbor to v. This information could be communicated through a local message exchange between u and v. This local message exchange is required only during planarization and so will not consume much overhead. It could also be piggybacked on Hello messages or location updates. Since planarization is required mainly after topology changes, the location update message could be a good candidate to include the extra information by adding an extension to provide the node's list of neighbors allowing other neighbors to apply the fix when required. In either case, whether this information is added as an extension or sent as a separate message, the additional overhead is small and we do not expect this change to have any noticeable effect on the operation.

This fix guarantees that the planar graph is always connected if the topology is connected, but it may add some extra cross-links. But, the probability of creating cross-links will still be low given the conditions mentioned earlier for cross-links. This fix also solves the disconnection problem caused by other reasons, such as obstacles and non-ideal radio ranges. We shall further investigate the efficacy of our proposed fix in the next section.

7. Simulations

In the previous section we have shown scenarios where errors in geographic routing happen due to location inaccuracy and the causes and conditions for these errors. We verified the *possibility* that these errors can happen in the crafted scenarios, but this alone does not show the *probability* of these errors happening in general topologies. In this section we use simulations to study the possibilities of errors happening in random topologies, in addition to their effects on performance.

The geographic routing protocol used in the simulations is GPSR [16] with *RNG* planarization.

[2] In addition, the disconnection problem seems to be also the most common problem caused by obstacles and irregular radio ranges as can be seen in the figures in Section 4.

[3] A similar fix was suggested (but not evaluated) in [15] to cope with obstacles.

Geographic routing is not the only application affected by localization inaccuracy; other systems and functions based on location information can also be affected. Accordingly, we study also the effect of localization inaccuracy on a geographic system, GHT [23], built on top of geographic routing. GHT uses face routing in another different way; to find the node closest to a certain geographic location. This will make routing failures more significant since face routing is invoked at every key insertion or lookup. It also introduces an additional kind of failure due to the inconsistent storage/retrieval that may result from inaccurate node locations.

We first run simulations for both GPSR and GHT at different densities with relatively small localization errors that we believe represent the current state-of-the-art localization systems. Then we evaluate the fix we introduced and show that they recover the most probable errors even with greater location errors.

7.1. Methodology and metrics

We are mainly interested in evaluating the effects of location inaccuracy on geographic routing independent of the MAC layer used and without concerning ourselves with other factors such as MAC collisions. Besides many of the MAC protocols proposed for sensor networks are not collision-based and since our focus is on sensor network scenarios with low traffic, the specific MAC protocol should not have a significant effect on the results. Thus, our simulations for GPSR and GHT consider only the routing behavior affected by the error conditions. We consider a static and stable network of 100 nodes having the same radio range of 80 m. We vary the density of the network by changing the space size, where the density is presented as the number of nodes per radio range. In each simulation run, nodes are placed at random locations in the topology and results are computed as the average of 1000 runs. We consider only topologies where the network is connected. The maximum localization error is presented as a fraction of the radio range. The estimated node location is picked uniformly from a random location around the node accurate position limited by the maximum localization error.

The main metric that concerns us in this study is the success rate of packet delivery since this represents the correctness of the protocol in the face of inaccuracy. We also evaluate the routing overhead to measure the effect on performance. In GPSR, a packet is sent from every node to every other node (this gives $n(n-1)$ routes among n nodes) and the success rate is computed as the percentage of packets delivered to the destination. In GHT, we assume 10 event types and for each type an event will happen by each node that will send a packet to the corresponding hash location. An access point sends a lookup for each type and the success rate is the percentage of events successfully retrieved from all events generated. For brevity, we only show the key results obtained.

7.2. Main results

In this section we present results for uniformly distributed random topologies with localization errors 1–10% of the radio range. We change the density from 5 nodes per range to 20 nodes per range and observe the success rate. Although, sensor networks deployed are expected to be of high density, the operational node density could be much less. Low-density networks are common either due to the environment or to improve the efficiency and power consumption. Several topology control techniques such as SPAN [6] and GAF [32] are proposed to save power by turning off nodes, which leads to small neighborhood size for each node and a sparse network. In addition, collisions at high-density networks increase the delay and overhead. The density is also not expected to be constant during the lifetime of the network; it will change due to node failures and power depletion.

Fig. 16(a) shows the success rate of GPSR. Even with relatively low location inaccuracy, the success rate is affected. At high densities (above 10) the success rate is above 99.5%, but all failures are persistent and non-recoverable, as mentioned earlier. At lower densities, the success rate decreases significantly. In Fig. 16(b), the success rate reduction in GHT is higher due to the face routing around the geographic hash location, which leads to more errors. In GHT errors happen also due to the inconsistency that can occur by storing at a node and retrieving from a different node. But our results show that the inconsistency errors are insignificant at these inaccuracy rates and start to emerge at higher inaccuracy. The overhead will increase slightly at these inaccuracy rates with larger increase at low density networks. The increase in overhead is due to the errors and the results also show that inaccuracy reduces greedy routing success and leads to more face routing. In GPSR the routing failures can lead

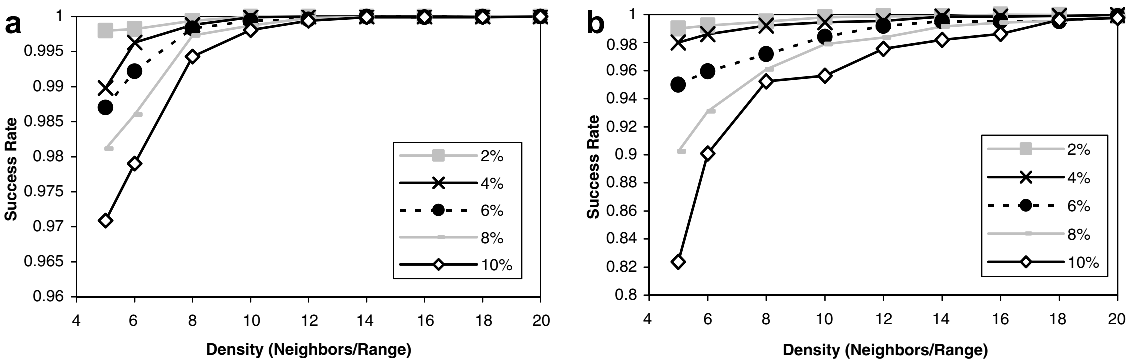

Fig. 16. The success rate of GPSR and GHT at different inaccuracy ranges (% of radio range) and densities. (a) GPSR (b) GHT.

to two error types: 1. Packet drops immediately after a node discovers that there is no route to the destination. This happens when a node sends a packet using face routing and receives the packet again without finding the destination. 2. Permanent loops before the TTL is exhausted and the packet is dropped. In GHT, the first error type of GPSR does not exist, since the packet is forwarded to a location and not to a specific destination, thus the node that initiates face routing and receives the packet back again with no closer nodes found considers itself the home node of the packet. Accordingly, in GHT the routing failure will lead only to permanent loops, in addition to the inconsistent storage/retrieval failure. Notice, that if some loop detection technique is used (which is inconvenient to implement in sensor networks due to resource constraints), previous permanent loops will become immediate packet drops.

From the figures it is clear that errors happen mainly at low densities, which gives us another indication that these errors are due to edge removal from the planar graph. To further validate this, we analyzed the simulation traces and classified routing errors into the 3 categories mentioned in Section 6.2: edge removal, cross links, and destination inaccuracy. Following are some of our main observations.

– Above 95% of the errors are due to edge removal only. About 5% of the errors happen as a combination of destination inaccuracy with edge removal. Less than 1% of the error paths contain cross-links.
– In GHT, more than 99% of the errors are due to edge removal and less than 1% of the errors show cross-links or inconsistency (note that inconsistency can also result from routing failures).

To evaluate the proposed fix, we ran the same simulations with the fix (of Section 6) added to GPSR and GHT. The success rate at all densities with a localization error range of 1–10% of the radio range is above 99.99% for GPSR and above 99% for GHT with almost all of the values 100%. This indicates that the simple fix added is good enough to fix almost all of the errors at least for the inaccuracy range of interest. This also shows that as we analyzed, planarization edge removal causes most of the errors (almost all of them in this range).

In addition to the nearly-prefect results the fix achieves at the low error ranges, at higher error ranges it also provides great improvements. The high error ranges could be due to large localization inaccuracies or due to faulty measurements. Fig. 17(a) shows the success rate of GPSR at high localization error ranges up to the whole radio range and Fig. 17(b) shows the success rate after applying the fix. The same is shown in Fig. 18 for GHT. For both GPSR and GHT our proposed modified (fixed) version of the protocol achieves over 97.5% success rate with up to 60% localization error, and over 85% success rate with up to 100% localization error, even with very low node density. The overhead also reduces significantly by adding the fix, which shows that the overhead of the fix itself is negligible comparable to the overhead of the problems it solves.

Furthermore, we conducted *ns*−2 [3] simulations with detailed models of the wireless MAC and physical layers. The general trends are similar, but the quantitative results are very sensitive to the traffic patterns and rates. The remarkable distinction related to our study is the effect of density. Though, higher density networks are more robust to location errors, they suffer from more collisions which affect their performance. Errors like permanent loops

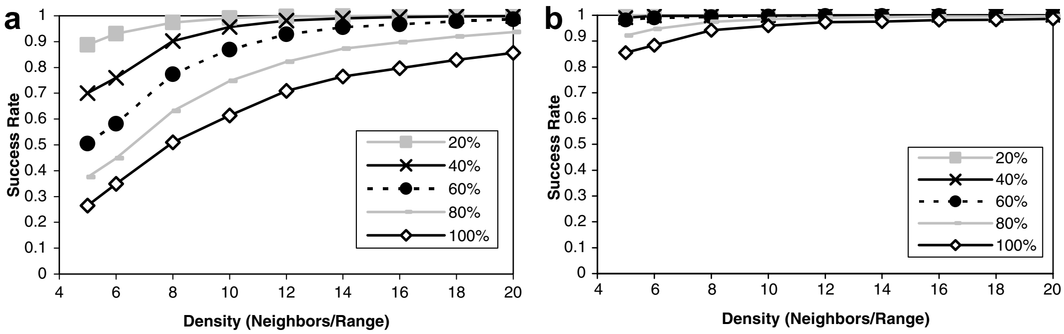

Fig. 17. The success rate of GPSR at high error ranges (% of radio range) *without* the fix and *with* the fix (a) without fix (b) with fix.

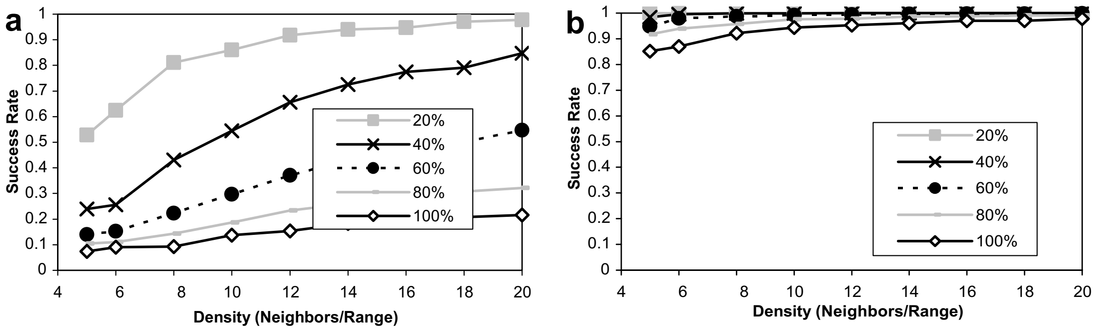

Fig. 18. The success rate of GHT at high error ranges (% of radio range) *without* the fix and *with* the fix (a) without fix (b) with fix.

have a severe effect on performance when collisions are considered. The extra overhead needed for the fix does not have any effect on the results, especially that the network is static.

7.3. Non-uniform distribution

In this section we study the effect of location errors and evaluate our fix when the nodes in the network are not uniformly distributed. Irregular distribution of nodes and the existence of gaps and obstacles cause more greedy routing failures and increase face routing. A large void area is set in the middle of the space, as shown in Fig. 19, with a side length equal to 80% of the total length. Fig. 20 shows the success rate without and with the fix. Without the fix, the success rate at high densities is lower at the non-uniform distribution compared to the uniform distribution. The fix still improves the success rate to above 80% with up to 100% localization errors.

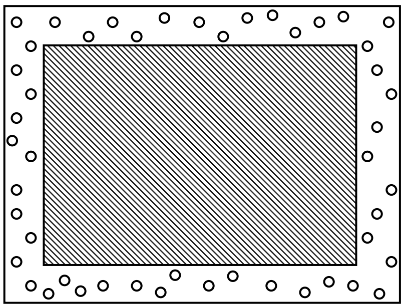

Fig. 19. A large void in the middle of space incurring more face routing.

7.4. Correlated error model

In this section we use a correlated error model instead of the uniform error model, in order to study the effect of location errors when the errors are correlated and to evaluate our fix in this environment. We assume that a reference exists at the centre of each quarter of the space as shown in Fig. 21. The error range of a node depends on how far it is from the closest reference. This reference could represent a beacon or a node that has accurate location in ad hoc localization systems. Fig. 22 shows the success rate using the correlated error model. Notice that the error range used in Fig. 22 is different than in the previous figures. It

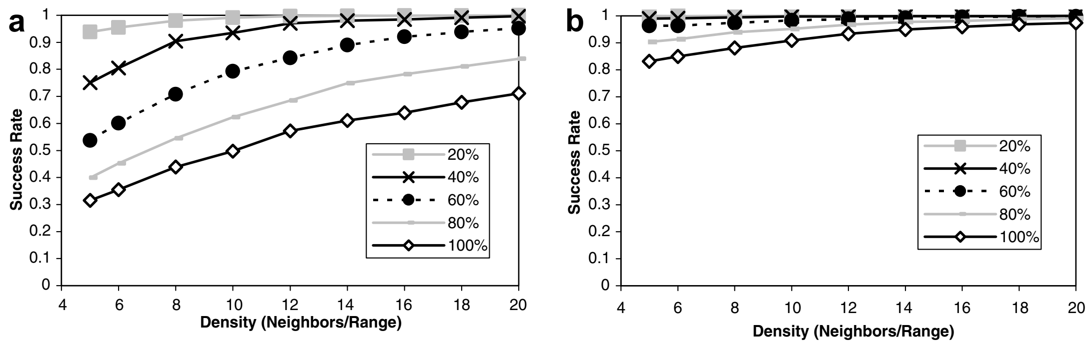

Fig. 20. The success rate of GPSR at a non-uniform distribution *without* the fix and *with* the fix (a) without fix (b) with fix.

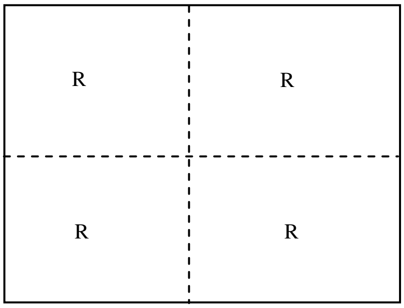

Fig. 21. The error range of a node depends on the distance between the node and the closest reference.

represents the maximum error range which is the error range of the farthest point from the reference. If the distance between a node and the closest reference is d, the maximum error range is e, and the maximum distance to a reference is $d_{\max}$ (in Fig. 21 this is the distance between the reference and the quarter corners), then the error range of that node is $e * d/d_{\max}$. Fig. 22(b) shows that the fix is efficient in solving most of the errors.

8. Conclusions

In this paper we presented a detailed systematic micro-level analysis of pathologies that occur in face routing-based geographic protocols. We have shown that conditions that violate the unit graph assumption such as location errors, obstacles, and irregular radio coverage cause planarization failures in the form of disconnections and cross-links and accordingly cause face routing failure. We then analyzed location errors in more detail and built scenarios for these errors. We adopt a novel approach in synthesizing the error scenarios; starting from the planarization algorithms we establish conditions for the errors. Based on this analysis, we presented a simple local fix that solves the most probable error which is graph disconnection. We further conducted simulation case studies (for GPSR and GHT) to quantify the effect of location errors on protocol performance and to validate the efficacy of our proposed modification at different error ranges, distributions and error models.

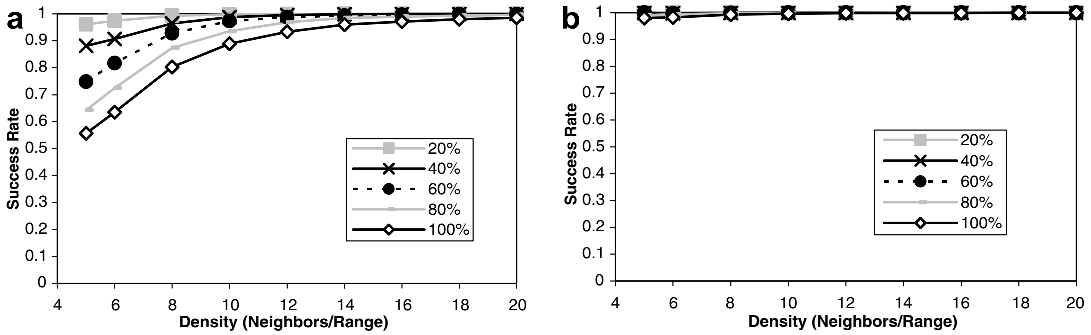

Fig. 22. The success rate of GPSR with a correlated error model *without* the fix and *with* the fix. The error range shown here is the maximum error range which is the error range of the farthest point from the reference (a) without fix (b) with fix.

We have shown by our micro-level analysis that failures in face routing happen for two reasons: disconnections in the planar graph or cross-links. The mutual-witness local fix presented solves one of these problems. It guarantees that the graph generated by the local planarization algorithms GG and RNG will be connected if the network is connected. But this local fix does not remove cross-links. Due to that, our simulation results have shown that this fix improves the success rate significantly since disconnections are more likely to cause failures, but it does not guarantee delivery because of cross-links. We have proven that local algorithms in general cannot solve all problems and cannot guarantee delivery under arbitrary connectivity. In order to guarantee delivery we need a non-local algorithm that can search or propagate information for an unlimited number of hops.

Acknowledgments

Seada and Helmy were partially supported by grants from NSF CAREER, Intel, and Pratt&Whitney. We would like to thank the anonymous reviewers for their feedback.

References

[1] L. Barrière, P. Fraigniaud, L. Narayanan, J. Opatrny, Robust position-based routing in wireless ad hoc networks with irregular transmission ranges, Wireless Communications and Mobile Computing 3 (2) (2003) 141–153.
[2] L. Blazevic, S. Giordano, J.-Y. Le Boudec, Self-organized terminode routing, Cluster Computing Journal 5 (2002).
[3] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, Y. Xu, H. Yu, Advances in network simulation, IEEE Computer (May) (2000).
[4] P. Bose, P. Morin, I. Stojmenovic, J. Urrutia, Routing with guaranteed delivery in ad hoc wireless networks, in: Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, DialM 1999.
[5] N. Bulusu, J. Heidemann, D. Estrin, T. Tran, Self-configuring localization systems: design and experimental evaluation, ACM Transactions on Embedded Computing Systems (TECS) (2003).
[6] B. Chen, K. Jamieson, H. Balakrishnan, R. Morris, Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks, in: ACM MOBICOM, 2001.
[7] E. Chávez, S. Dobrev, E. Kranakis, J. Opatrny, L. Stacho, J. Urrutia, Route discovery with constant memory in oriented planar geometric networks, in: ALGOSENSORS, 2004, pp. 147–156.
[8] G. G. Finn, Routing and addressing problems in large metropolitan-scale internetworks, Technical Report ISI/RR-87-180, Information Sciences Institute, March 1987.
[9] K. Gabriel, R. Sokal, A new statistical approach to geographic variation analysis, Systematic Zoology (1969).
[10] B. Greenstein, D. Estrin, R. Govindan, S. Ratnasamy, S. Shenker, DIFS: a distributed index for features in sensor networks, in: IEEE SNPA, 2003.
[11] D. Goldenberg, A. Krishnamurthy, W. Maness, Y.R. Yang, A. Young, A. Savvides, Network localization in partially localizable networks, in: IEEE INFCOM, 2005.
[12] T. He, C. Huang, B. Blum, J. Stankovic, T. Abdelzaher, Range-free localization schemes for large scale sensor networks, in: ACM MOBICOM, 2003.
[13] J. Hightower, G. Borriello, Location systems for ubiquitous computing, IEEE Computer (August) (2001).
[14] R. Jain, A. Puri, R. Sengupta, Geographical routing using partial information for wireless ad hoc networks, IEEE Personal Communications (February) (2001).
[15] B. Karp, Challenges in geographic routing: sparse networks, obstacles, and traffic provisioning, Slides Presented at the DIMACS workshop on Pervasive Networking, May 2001. (www.icir.org/bkarp/gpsr/gpsr.html).
[16] B. Karp, H.T. Kung, GPSR: greedy perimeter stateless routing for wireless networks, in: ACM MOBICOM, 2000.
[17] Y.-J. Kim, R. Govindan, B. Karp, S. Shenker, Geographic routing made practical, in: Proceedings of the USENIX Symposium on Networked Systems Design and Implementation, May 2005.
[18] Y.-J. Kim, R. Govindan, B. Karp, S. Shenker, On the Pitfalls of Geographic Routing, in: Proceedings of the 3rd International Workshop on DIALM-Principles of Mobile Computing, September 2005.
[19] E. Kranakis, H. Singh, J. Urrutia, Compass routing on geometric networks, in: Proceedings of the 11th Canadian Conference on Computational Geometry, August 1999.
[20] F. Kuhn, R. Wattenhofer, A. Zollinger, Worst-case optimal and average-case efficient geometric ad-hoc routing, in: ACM Mobihoc, 2003.
[21] D. Moore, J. Leonard, D. Rus, S. Teller, Robust distributed network localization with noisy range measurements, in: SenSys, 2004, pp. 50–61.
[22] N. B. Priyantha, A. Chakraborty, H. Balakrishnan, The cricket location-support system, in: ACM MOBICOM, 2000.
[23] S. Ratnasamy, B. Karp, L. Yin, F. Yu, D. Estrin, R. Govindan, S. Shenker, GHT: A geographic hash table for data-centric storage, in: ACM WSNA, 2002.
[24] Savvides, C.-C. Han, M. B. Srivastava, Dynamic fine-grain localization in ad-hoc networks of sensors, in: ACM MOBICOM, 2001.
[25] K. Seada, A. Helmy, Geographic protocols in sensor networks, in: Encyclopedia of Sensors, American Scientific Publishers (ASP), 2006.
[26] K. Seada, A. Helmy, R. Govindan, On the effect of localization errors on geographic face routing in sensor networks, USC Technical Report, June 2003. in: IEEE/ACM 3rd International Symposium on Information Processing in Sensor Networks (IPSN), Berkeley, CA, April 2004.
[27] K. Seada, M. Zuniga, A. Helmy, B. Krishnamachari, Energy-efficient forwarding strategies for geographic routing in lossy wireless sensor networks, in: ACM SenSys, 2004.
[28] R. Shah, A. Wolisz, J. Rabaey, On the performance of geographical routing in the presence of localization errors, IEEE International Conference on Communications (2005).

[29] S. Subramanian, S. Shakkottai, Geographic routing with limited information in sensor networks, in: ACM IPSN, April 2005.
[30] G. Toussaint, The relative neighborhood graph of a finite planar set, Pattern Recognition 12 (4) (1980) 261–268.
[31] A. Ward, A. Jones, A. Hopper, A new location technique for the active office, IEEE Personal Communications (October) (1997).
[32] Y. Xu, J. Heidemann, D. Estrin, Geography-informed energy conservation for ad-hoc routing, in: ACM MOBICOM, 2001.

Karim Seada received his Ph.D. in Electrical Engineering and M.S. in Computer Science from the University of Southern California, and his M.S. and B.S. in Computer Engineering from Cairo University. He is currently a member of Nokia Research Center in Palo Alto, California, working on researching innovative services for wireless networks. His research interests include wireless mesh, ad hoc, and sensor networks, in addition to geographic services and location-based protocols. He held previous internships at Intel and Dust Networks.

Ahmed Helmy received his Ph.D. in Computer Science (1999), MS in Electrical Engineering (1995) from the University of Southern California, MS Eng Math (1994) and BS in Electronics and Communications Engineering (1992) from Cairo University, Egypt. He is now Associate Professor at the Computer and Information Science and Engineering department at the University of Florida. From 1999 to 2006 he was an Assistant Professor of Electrical Engineering at the University of Southern California. In 2002, he received the National Science Foundation (NSF) CAREER Award. In 2000 he received the USC Zumberge Research Award, and in 2002 he received the best paper award from the IEEE/IFIP International Conference on Management of Multimedia Networks and Services (MMNS). He is co-leading the STRESS and ACQUIRE NSF-funded projects. He is the founder and director of the wireless networking laboratory at UFL (and previously at USC). His current research interests lie in the areas of protocol design and analysis for mobile ad hoc and sensor networks, mobility modeling, design and testing of multicast protocols, IP micro-mobility, and network simulation.

Ramesh Govindan received his B. Tech. degree from the Indian Institute of Technology at Madras, and his M.S. and Ph.D. degrees from the University of California at Berkeley. He is an Associate Professor in the Computer Science Department at the University of Southern California. His research interests include scalable routing in internetworks, and wireless sensor networks.