

On the Effect of Localization Errors on Geographic Face Routing in Sensor Networks

Karim Seada

Electrical Engineering Department
University of Southern California

seada@usc.edu

Ahmed Helmy

Electrical Engineering Department
University of Southern California

helmy@usc.edu

Ramesh Govindan

Computer Science Department
University of Southern California

ramesh@usc.edu

ABSTRACT

In the absence of location errors, geographic routing - using a combination of greedy forwarding and face routing - has been shown to work correctly and efficiently. The effects of location errors on geographic routing have not been studied before. In this work we provide a detailed analysis of the effects of location errors on the correctness and performance of geographic routing in static sensor networks. First, we perform a micro-level behavioral analysis to identify the possible protocol error scenarios and their conditions and bounds. Then, we present results from an extensive simulation study of GPSR and GHT to quantify the performance degradation due to location errors. Our results show that even small location errors (of 10% of the radio range or less) can in fact lead to incorrect (non-recoverable) geographic routing with noticeable performance degradation. We then introduce a simple modification for face routing that eliminates probable errors and leads to near perfect performance.

Categories and Subject Descriptors

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

General Terms

Algorithms, Performance, Design, Reliability, Verification.

Keywords

Geographic Routing, Face Routing, Localization Inaccuracy, Location Errors, Wireless Sensor Networks, Data-Centric Storage.

1. INTRODUCTION

Geographic routing protocols ([11], [13]) 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 nodes need not maintain per-destination information, and only neighbor location information is needed to route packets. For these reasons, geographic routing is becoming the protocol of choice for many emerging applications in sensor networks, such as data-centric storage [17] and distributed indexing [7]. 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 all geographic routing protocols till date have assumed the availability of accurate location information. In practice (in systems that either rely entirely on GPS, or infer location using ad-hoc localization systems [9]), however, location measurement is often noisy and incurs some error. For example, many state-of-the-art techniques usually incur around 10% (of the radio range) or more in localization error. To our knowledge there has been no previous detailed study on the effects of localization errors on the correctness and performance of geographic routing. This paper attempts to fill that void.

In this paper, we first analyze the pathologies that can arise in face-routing based geographic routing protocols, in the presence of errors in node location. Our methodology for this analysis is novel: using an elaborate, micro-level analysis of face routing protocols, we provide detailed scenarios in which protocol correctness is violated when the location of a node is in error.

We then perform extensive simulations to evaluate and quantify the effects of localization errors on two prominent protocols that use face routing: GPSR [11] and GHT [17]. Our study shows that realistic localization 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 that eliminates the most likely protocol errors and we evaluate the efficacy of our fix. Our simulations show near perfect performance for our modified geographic routing (for GPSR and GHT) even in the presence of significant localization errors.

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 detailed micro-level analysis and the error scenarios. Section 5 discusses the fix

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.

IPSN'04, April 26–27, 2004, Berkeley, California, USA.
Copyright 2004 ACM 1-58113-846-6/04/0004...\$5.00.

and its rationale. Section 6 contains the simulation results. Finally, conclusions are presented in Section 7.

2. RELATED WORK AND BACKGROUND

Early work in geographic routing considered only greedy forwarding [5] 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 [12] 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.* [3] 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.

GPSR [11] 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. GOAFR [13] is a recently proposed protocol that also combines greedy forwarding with face routing and achieves worst-case optimality of the path length by using limited elliptic regions for face routing and recursively increasing the ellipse size until finding a close-to-optimal path.

More recently, sensor network researchers have proposed several systems that build upon GPSR in order to support data-centric storage. GHT [17] 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 GPSR for geographic routing to the hash location. 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 [7]. DIFS provides a distributed index for efficient index construction and range searches in sensor networks.

In [8], 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. In this paper, we provide a detailed analysis for the effect of location errors on complete geographic routing protocols consisting of greedy forwarding coupled with face routing. We show more concrete results and provide the reasons for errors.

To our knowledge, no prior work has analyzed the effect of location errors either on geographic face routing protocols, or on data-centric storage systems built on top of these protocols.

Many localization systems have been proposed in the literature: GPS, infrastructure-based localization systems [20][15], and ad-

hoc localization systems [2][18]. We do not discuss these systems in detail; the interested reader is referred to an extensive survey of localization by Hightower *et al.* [9]. 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.

Recently, there have been proposals [16][14] for geographic routing without location information. The approach is to assign logical coordinates to each node and then use greedy forwarding over the logical coordinates. In order to build the logical space, a relatively large setup overhead is required in the form of several global floods or iterations of coordinates' relaxation. Extra overhead is incurred to detect and propagate changes in the logical coordinates due to topology changes. Although it has some benefits, routing *without* location information is not as efficient and scalable as routing *with* location information. Perimeter routing also cannot be used in these schemes, and in [16] there is no solution provided for dead-ends that guarantees delivery. In addition, for data-centric storage, these schemes do not provide robust replication and consistency.

3. MODEL AND ASSUMPTIONS

Before we analyze the impact of location errors 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 consistent view of the network. Admittedly, these are idealized assumptions that do not hold in practice. As we illustrate in this paper, positional inaccuracies can result in routing pathologies *even under the rest of idealized assumptions*.

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 [3][11][13].

Our study focuses on the effect of inaccurate localization errors on geographic routing. Thus, we assume a static and stable network (no mobility and no failures) without obstacles and with nodes having accurate and symmetric radio ranges. We also assume that nodes have consistent location information about other nodes, which means that a node estimates its location and announces it, and all nodes observe the same estimated location for that node.

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 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. The error scenarios apply to both *RNG* and *GG* planarization algorithms, with locations adjusted to satisfy the failures conditions.

Figure 4 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, Figure 4(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 . In the estimated topology, Figure 4(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 Figure 5, the opposite happens. Edge (F_1, F_2) is removed by planarization in the accurate topology, Figure 5(a), because node E is a witness, allowing face routing to deliver the packet from S to D . In the estimated topology, in Figure 5(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. 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 Figure 6, 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 by the local planarization algorithm. Since face routing assumes and requires planar graphs, cross-links cause route failure.

Figure 7 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.

In Section 5, we use GPSR [11] and GHT [17] simulations to verify that these errors do actually occur and to quantify the probability of their occurrence. The problems uncovered by our scenarios also exist in other protocols using face routing. For example, Compass Routing II [12] uses a version of face routing where packets are forwarded along the faces intersecting the line connecting the source and the destination. GOAFR [13] is a newly proposed protocol that also suffers from the same problems with inaccuracy. GOAFR is similar to GPSR in combining greedy forwarding with face routing on a planarized graph. It differs in bounding the face by an ellipse and face routing within that ellipse, which is a performance enhancement and does not affect the correctness failures incurred in our scenarios. Another difference is that GOAFR switches faces at the face point closest to the destination. This will also not affect the failures in the presented scenarios.




Figure 4: Disconnection due to incorrect edge removal by planarization




Figure 5: Permanent loop due to planarization failure in not removing edge

Obstacles and Non-Ideal Radio Range

Although our focus in this paper is on errors due to inaccurate geographic locations, similar errors could also happen for other reasons such as obstacles and non-ideal radio ranges. In many environments obstacles are natural and can cause two proximal nodes to be physically disconnected. This also violates the assumption of the unit graph and causes a planarization failure. For example, in Figure 4(b), assume nodes have accurate

locations but there is an obstacle between E and F_1 such that S removes the edge (S, F_1) from the planar graph while there is no other path through E . This causes an error similar to that resulting from inaccuracy. Another example also in Figure 4(b), assume that the radio range of F_1 is not ideal (by ideal here we mean the common assumption of perfect spherical, symmetric and accurate radio range) such that E does not hear F_1 . The removal of (S, F_1) from the planar graph causes disconnection. Furthermore, any other cause of *information hiding* such that one of the nodes cannot get the information to construct a correct planar graph, will lead to the failure of face routing.




Figure 6: Cross links causing face routing failure




Figure 7: Destination inaccuracy causing failure in reaching the destination. The route in this figure assumes that F_2 is the closest node to D^* . The packet is forwarded greedily to F_2 , and then face routing is used around the perimeter

4.3 Error Analysis

In this section we analyze the protocol components to show the completeness of our error scenarios under the current model and assumptions.

4.3.1 Unit Graph Assumption

A main assumption violated by location inaccuracy is that the topology is a unit graph, which means that two nodes are connected if and only if the Euclidean distance between them is less than the radio range. More formally,

u and v are connected iff $d(u, v) < R$, where R is the radio range
if $v \in N \Rightarrow d(u, v) < R$, where N is the neighbor list of u

if $d(u, v) < R \Rightarrow v \in N$

Inaccuracy can violate this assumption in two ways:

$v \in N$ and $d(u^{\wedge}, v^{\wedge}) > R$ (u^{\wedge} and v^{\wedge} are the estimated locations)
 $d(u^{\wedge}, v^{\wedge}) < R$ and $v \notin N$

The planarization is based on the assumption that the topology is a unit graph, so violation of this assumption causes wrong decisions by the planarization algorithm. It could also cause cross-links that cannot be resolved by a *local* planarization algorithm.

4.3.2 Planarization Algorithm

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 add an edge that should be removed.

In *RNG* (Figure 1) an error will happen when
 $\text{decision}\{d(u, v) > \max[d(u, w), d(w, v)]\} \neq$
 $\text{decision}\{d(u^{\wedge}, v^{\wedge}) > \max[d(u^{\wedge}, w^{\wedge}), d(w^{\wedge}, v^{\wedge})]\}$

While in *GG* (Figure 2) the error will happen when
 $\text{decision}\{d(c, w) < d(c, u)\} \neq$
 $\text{decision}\{d(c^{\wedge}, w^{\wedge}) < d(c^{\wedge}, u^{\wedge})\}$

4.3.3 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 Figure 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:

- a disconnection in the planar graph (Scenario 1)
- or a cross-link (Scenarios 2 and 3)

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:

- the destination is not at the estimated location (Scenario 4)
- nodes at the estimated location face have no access to the destination

In summary, we have shown from first principles that the scenarios described above *completely* characterize the kinds of pathologies that can arise in face routing.

4.4 Conditions and Bounds

This section more precisely describes the geometric conditions under which face routing failures can arise as well as bounds for the planarization errors that cause face routing failure. As we have seen above, the different types of errors are: an edge removed that should not have been removed, edge not removed that should have been removed, and cross-links. The bounds presented are assuming only a single node is inaccurate, but the conditions are general for multiple node inaccuracy.

4.4.1 Error: Excessive Edge Removal

Excessive edge removal occurs when an edge is removed from planar graph that should not be removed, which may lead to disconnection.

Conditions: Assume $u=(x,y)$ is the exact location and $u^=(x^,y^)$ is the estimated edge location. Following are the necessary condition for removing edge (u,v) that may cause disconnection:

\exists no w such that $d(u,v) > \max[d(u,w), d(w,v)]$
 $\& \exists w$ such that $d(u^,v^) > \max[d(u^,w^), d(w^,v^)]$
 (where w is a neighbor of u , but not of v)

Or in a different form

$d(u,v) < R, d(u,w) < R,$
 $d(w,v) > R$ (this condition is for disconnection between u and v to be possible, otherwise an edge (u,v) will be removed and another (w,v) added)
 $d(u,v) < \max[d(u,w), d(w,v)] \forall w \in N$
 $d(u^,v^) > \max[d(u^,w^), d(w^,v^)]$ for some $w \in N$

Bounds: The question we want to answer here is the following: for a given u, v , and w , what is the minimum $\Delta = |w-w^|$ to cause this error (bounds on Δ)?

RNG planarization:

RNG intersection equation:

$$u \text{ circle: } (x-x_u)^2 + (y-y_u)^2 < d(u,v)^2$$

$$v \text{ circle: } (x-x_v)^2 + (y-y_v)^2 < d(u,v)^2$$

without loss of generality assume $x_u = 0, y_u = 0,$
 $x_v = r = d(u,v), y_v = 0$

Equation of the intersection curve:

$$y^2 = \frac{2xr - x^2}{r^2 - x^2} \quad 0 < x < r/2$$

$$y^2 = \frac{2xr - x^2}{r^2 - x^2} \quad r/2 < x < r$$

$\Delta >$ distance between w and closest point on intersection curve

$\Delta <$ distance between w and furthest point on intersection curve

Approximately:

$$\Delta > d(v,w) - d(u,v) \quad (\text{Figure 8})$$

For Gabriel Graph planarization (Figure 9):

$$d(c,w) - d(u,v)/2 < \Delta < d(c,w) + d(u,v)/2$$

For multiple neighbors, any neighbor having an estimated position inside the intersection causes the edge to be removed.




Figure 8: RNG bound




Figure 9: GG bound

4.4.2 Error: Insufficient Edge Removal

Insufficient edge removal occurs when an edge is not removed from planar graph that should have been removed, which may cause loops.

Conditions: Following are the conditions that cause an edge not to be removed from the planar graph when it should have been removed:

$\exists w$ such that $d(u,v) > \max[d(u,w), d(w,v)]$ &
 \exists no w such that $d(u^,v^) > \max[d(u^,w^), d(w^,v^)]$
 (where w is a neighbor of u and v)

Or in a different form

$d(u,v) < R, d(u,w) < R, d(w,v) < R$
 $d(u,v) > \max[d(u,w), d(w,v)]$ for some $w \in N$
 $d(u^,v^) < \max[d(u^,w^), d(w^,v^)] \forall w \in N$

Bounds: The bounds for RNG and GG are:

(a) RNG planarization:

$\Delta >$ distance between w and closest point on intersection curve

(b) Gabriel Graph planarization:

$\Delta > d(u,v)/2 - d(c,w)$

For multiple neighbors, all neighbors in the intersection should have estimated positions out of the intersection for the edge not to be removed.

4.4.3 Creating Cross-Links

The planarization algorithm (RNG or GG) is supposed to remove cross-links from a unit graph and form a planar graph. Due to location errors, cross-links may not be removed. The conditions for cross-links are explained next.




Figure 10: Case 1 or 2




Figure 11: Case 3

Conditions: The first condition is that x is not connected to u or v , i.e., $(d(x,u) > R \text{ and } d(x,v) > R)$. (if it is connected to any of them, the cross link will be removed by planarization).

The remaining conditions depend on the state of w . There are 3 different cases depending on the accurate position of w with regard to u and v (see Figure 10 and Figure 11).

1. w is not connected to u or v
2. w is connected to u or v , and (x,w) does not cross (u,v)
3. (x,w) crosses (u,v)

Case 1: w is not connected to u or v

Conditions:

(x, w) does not cross (u, v)
 (x^{\wedge}, w^{\wedge}) crosses (u^{\wedge}, v^{\wedge})

Bounds:

$\Delta >$ distance from w to (u, v)

Case 2: w is connected to u or v , and (x, w) does not cross (u, v)

Conditions:

(x, w) does not cross (u, v)
 (x^{\wedge}, w^{\wedge}) crosses (u^{\wedge}, v^{\wedge})
 w^{\wedge} is not in (u^{\wedge}, v^{\wedge}) planar intersection

Bounds:

$\Delta >$ distance from w to (u, v) + distance to closest point on (u, v) intersection curve

Case 3: (x, w) crosses (u, v)

Conditions:

(x, w) crosses (u, v) , w is in (u, v) planar intersection
 (x^{\wedge}, w^{\wedge}) crosses (u^{\wedge}, v^{\wedge})
 w^{\wedge} is not in (u^{\wedge}, v^{\wedge}) planar intersection

Bounds:

$\Delta >$ distance from w to closest point on (u, v) intersection curve

This case is also an insufficient edge removal error.

5. DISCUSSION AND FIXES

Based on the understanding developed in the micro-level analysis section, we now propose fixes for some of the above problems and explain our rationale for choosing those fixes. Not all the problems are equally likely to occur. Based on their cause we reason about the probability of occurrence and provide fixes for what we believe to be the most likely to occur. The fixes are further evaluated, quantitatively, in the simulation section. In addition, the simulation section results confirm many of the points we argue about here. From the micro-level analysis, Section 4, we can categorize the problems that cause face routing failure into three main categories: 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 and provide fixes that solve most of them.

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 (see Section 4.4.1). This is a reasonably possible error in sparse networks even with low inaccuracy. The second problem of cross-links (Section 4.4.3) with or without a corresponding 'insufficient edge removal' error; from figure 10 node w needs to: (case 1) cross edge (u, v) if it is not connected to any of them or (case 2) cross edge (u, v) and get out of their planarization intersection if it is connected to any of them (case 2). This is very unlikely under the considered (10%) inaccuracy range, since in

case 1, if w is not connected to any of them, that means it must be at a distance higher than the radio range. Hence, with location errors on the order of 10% of the radio range it is highly improbable that a cross edge will be created. And in case 2, w needs both to cross the edge and pass the intersection, also an unlikely event. In case 3 (Figure 11), node w needs to get out of the planarization intersection, under the condition that x is not connected to u or v and (x, w) crossing (u, v) , which has also a low probability. In addition, for the edges (u, v) and (x, w) both to be added to the planar graph all nodes inside the corresponding intersections need to get out. 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 performance¹. From the planarization algorithm, an edge is removed from the planar graph when a witness is seen by a node (e.g., in Figure 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)².

This fix requires modification to the planarization algorithm as follows. Before removing an edge (u, v) , a node, u , sends a message to its neighbor, v , to inquire whether v sees the witness w . Node u must *not* remove the edge until and unless it gets a reply from v , indicating that v indeed sees w .

This message exchange is local between neighbors and it is required only during planarization and so will not consume much overhead. This fix guarantees that the planar graph is always connected if the topology is connected, but its drawback is that 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.

6. SIMULATIONS

In the micro-level analysis 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. In

¹ In addition, the disconnection problem is the problem also caused by obstacles and non-ideal radio ranges.

² A similar fix was suggested (but not evaluated) in [10] to cope with obstacles.

arguing our case, we take the following position: *Since these errors are correctness errors that lead to non-recoverable persistent routing failures to reachable destinations, even small percentages of these errors are not normally acceptable in static and stable networks.*

The geographic routing protocol used in the simulations is GPSR [11] with RNG planarization³. 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 [17], 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 fixes we introduced and show that they recover the most probable errors even with greater location errors.

6.1 Methodology and Metrics

We are mainly interested in evaluating the effects of location inaccuracy on geographic routing without any interference from other layers such as MAC collisions or physical layers effects. Thus, our simulations for GPSR and GHT consider only the routing behavior in an ideal wireless environment. We consider a static and stable network of 100 nodes having the same radio range of 80 meters. 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. 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

³ We ran the simulations also for GG and the effects of inaccuracy are still significant although with an improved success rate compared to RNG because GG planarization is denser and has less edge removals. We show the results for RNG, which was also the algorithm of choice in GPSR evaluations [11], because its lower density of connectivity reduces the MAC layer contentions and generally should improve performance. Existing MAC layer protocols (e.g. 802.11) show drastically reduced efficiency as the number of sending nodes in range increases.

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 in our study.

6.2 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 deployment are expected to be of high density, the operational node density could be much less. Low dense networks are common either due to the environment or to improve the efficiency and power consumption. Several topology control techniques such as SPAN [4] and GAF [21] 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 dense 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.

Figure 12(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 Figure 12(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 starts 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 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 5: 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.
- Above 70% of the errors lead to permanent loops.
- 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).




Figure 12: The success rate of GPSR and GHT at different inaccuracy ranges (% of radio range) and densities

6.3 Simulations with Fixes

To evaluate the proposed fix, we run the same simulations with the fix (of Section 5) 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 expected, planarization edge removal causes most of the errors (almost all of them in this range).

6.4 Sensitivity Analysis

In addition to the near-perfect 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. Figure 13(a) shows the success rate of GPSR at high localization error ranges up to the whole radio range and Figure 13(b) shows the success rate after applying the fix. The same is shown in Figure 14

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.

The error classification at this range shows the following:

- Edge removal error percentage increases at lower density.
- Destination inaccuracy and cross-links percentages increase with the error range. Destination inaccuracy errors are more significant than cross-links.

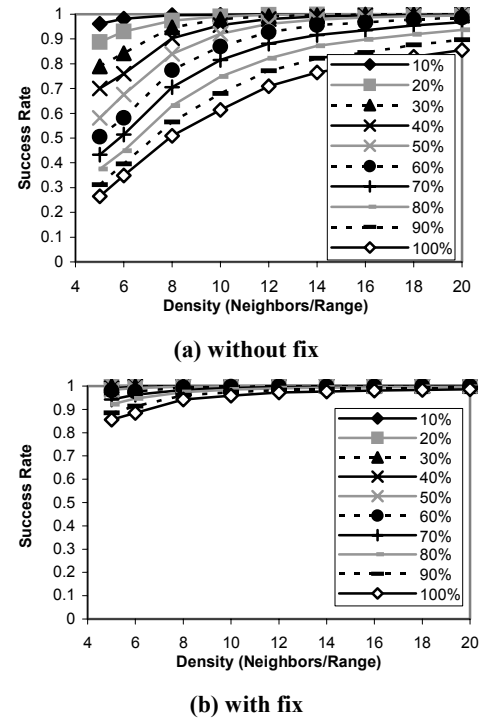


Figure 13: The success rate of GPSR at high error ranges (% of radio range) without the fix and with the fix

Furthermore, we conducted *ns-2* [1] simulations with detailed models of the wireless MAC and physical layers. The general trends are similar, but the quantitative results are so 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 have severe effect on performance when collisions are considered.

7. CONCLUSIONS

In this paper we presented the first detailed micro-level analysis of pathologies for geographic face-based routing protocols, in the presence of location errors in static sensor networks. We adopt a novel approach in synthesizing the error scenarios; starting from the planarization algorithms we establish completeness conditions and bounds for errors. Our analysis identified *information hiding* as one of the main causes for incorrect (non-recoverable)

behavior. We then provided a simple fix for this error based on information sharing among the nodes during planarization. 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. Our results show that even for realistic and relatively small location errors, the effects of location errors are noticeable. For example, in GHT more than 10% of sensor network events storage can fail in the presence of 10% location error. Furthermore, the validation of our proposed fix shows very promising results. Our modified GPSR and GHT versions achieve near-perfect performance even in the presence of significant localization errors.

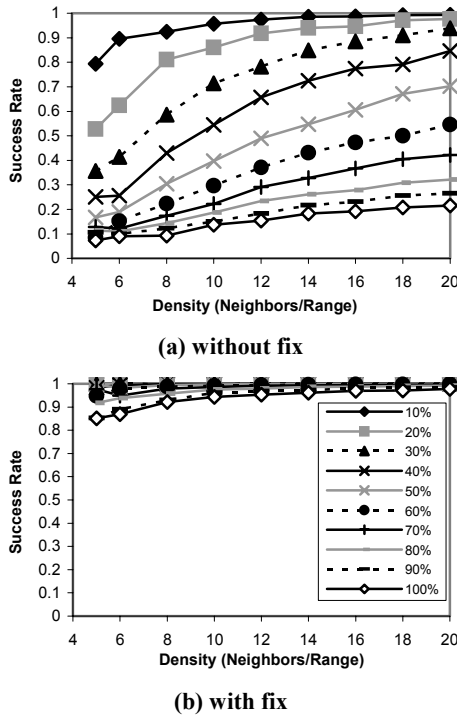


Figure 14: The success rate of GHT at high error ranges (% of radio range) without the fix and with the fix

8. ACKNOWLEDGMENTS

Seada and Helmy were partially supported by grants from NSF CAREER, Intel, and Pratt&Whitney. We would like to thank Brad Karp and the anonymous reviewers for their useful feedback. We like to thank also Sylvia Ratnasamy for providing the GHT simulation code.

9. REFERENCES

- [1] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, Y. Xu, and H. Yu. "Advances in Network Simulation." *IEEE Computer*, May 2000.
- [2] N. Bulusu, J. Heidemann, D. Estrin, and T. Tran. "Self-configuring Localization Systems: Design and Experimental Evaluation". *ACM Transactions on Embedded Computing Systems (TECS)* 2003.
- [3] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. "Routing with Guaranteed Delivery in Ad Hoc Wireless Networks".

Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DialM 1999).

- [4] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. "Span: An Energy-Efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks". *ACM MOBICOM* 2001.
- [5] G. G. Finn. "Routing and Addressing Problems in Large Metropolitan-Scale Internetworks". *Tech. Rep. ISI/RR-87-180*, Information Sciences Institute, Mar. 1987.
- [6] K. Gabriel and R. Sokal. "A New Statistical Approach to Geographic Variation Analysis". *Systematic Zoology*, 1969.
- [7] B. Greenstein, D. Estrin, R. Govindan, S. Ratnasamy, and S. Shenker. "DIFS: A Distributed Index for Features in Sensor Networks". *IEEE SNPA* 2003.
- [8] T. He, C. Huang, B. Blum, J. Stankovic, and T. Abdelzaher. "Range-Free Localization Schemes for Large Scale Sensor Networks". *ACM MOBICOM* 2003.
- [9] J. Hightower and G. Borriello. "Location Systems for Ubiquitous Computing". *IEEE Computer*, Aug. 2001.
- [10] 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)
- [11] B. Karp and H.T. Kung. "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks". *ACM MOBICOM* 2000.
- [12] E. Kranakis, H. Singh, and J. Urrutia. "Compass Routing on Geometric Networks". In *Proc. 11th Canadian Conference on Computational Geometry*, August 1999.
- [13] F. Kuhn, R. Wattenhofer, and A. Zollinger. "Worst-Case Optimal and Average-Case Efficient Geometric Ad-Hoc Routing". *ACM Mobihoc* 2003.
- [14] James Newsome and Dawn Song. "GEM: Graph Embedding for Routing and Data-Centric Storage in Sensor Networks without Geographic Information". *ACM Sensys*, November 2003.
- [15] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. "The Cricket Location-Support System". *ACM MOBICOM* 2000.
- [16] Ananth Rao, Sylvia Ratnasamy, Christos Papadimitriou, Scott Shenker, and Ion Stoica. "Geographic Routing without Location Information". *ACM MOBICOM*, September 2003.
- [17] S. Ratnasamy, B. Karp, L. Yin, F. Yu, D. Estrin, R. Govindan, and S. Shenker. "GHT: A Geographic Hash Table for Data-Centric Storage". *ACM WSNA* 2002.
- [18] A. Savvides, C.-C. Han, and M. B. Srivastava. "Dynamic Fine-Grain Localization in Ad-Hoc Networks of Sensors". *ACM MOBICOM* 2001.
- [19] G. Toussaint. "The Relative Neighborhood Graph of a Finite Planar Set". *Pattern Recognition* 12, 4 (1980), 261–268.
- [20] A. Ward, A. Jones, and A. Hopper. "A New Location Technique for the Active Office". *IEEE Personal Communications*, October 1997.
- [21] Y. Xu, J. Heidemann, and D. Estrin. "Geography-informed Energy Conservation for Ad-hoc Routing". *ACM MOBICOM* 2001.