This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2009 proceedings

On the benefits of nondeterminism in location-based forwarding

Israat Tanzeena Haque, Ioanis Nikolaidis, and Pawel Gburzynski

Department of Computing Science
University of Alberta
Edmonton, Alberta, CANADA T6G 2E8
Email: [israat,yannis,pawel]@cs.ualberta.ca

***Abstract*—We investigate the performance of location-based wireless ad-hoc forwarding schemes in generic network configurations, including regular grids as well as random networks. Through extensive simulation experiments in a realistic wireless propagation model, we arrive at the conclusion that winning forwarding strategies must be randomized, even if the population of next hop neighbors considered for the random selection includes nodes that would be obvious bad choices, if viewed as deterministic candidates. Inspired by this observation, we propose two randomized forwarding strategies and demonstrate their advantages over deterministic schemes. Notably, their advantage becomes more pronounced when traffic patterns are biased, where the primary advantage of nondeterminism is in circumventing hot spots.**

I. Introduction

We study the impact of randomization on the performance of location based routing schemes. Consider a simple (deterministic) *greedy* approach whereby a packet is always forwarded to the neighboring node that offers the shortest residual distance to the destination. Such a scheme exhibits good (and practically unbeatable) performance on random well-connected networks and under uniform load scenarios, which, however, tends to deteriorate in biased network topologies and/or under biased traffic patterns—because of hot spots. While a randomized protocol may be able to avoid hot spots (by statistically diversifying the routes), it may incur "deflections" resulting in longer routes, which will tend to reduce the network throughput. This in a nutshell explains the trade-offs inherent in the two approaches.

Our primary objective is to investigate whether randomization strategies, representing the wireless equivalents of what is frequently termed "deflection" [1], [8] in wired networks, are effective also in ad-hoc wireless networks. We consider deflection as being synergistic with location-based and geometric routing. The latter reduces routing information to a trivially small amount, while the former provides for plausible, albeit not necessarily optimal, strategies for forwarding.

II. Related Work

Numerous location-based routing schemes have been covered in the literature, including variations on the greedy theme [7], [11], [15], attempts at guaranteed packet delivery [3], [10], and MAC-assist cross-layer solutions [4], [16]. In this paper we focus on the first class of protocols, which will give us the purest framework for assessing the relevance of randomization. Protocols in that category strive to "make progress" in every hop, understood as reducing some measure of separation of the packet from its destination.

With *Greedy Routing* [7], the routing node selects the next-hop neighbor as the one with the shortest distance to the destination. Although the scheme works statistically well in reasonably dense networks, it suffers from the problem of *local maxima* (i.e., cul-de-sac's). Two variants of Greedy routing that circumvent this problem are GPSR (Greedy Perimeter Stateless Routing) and GFG (Greedy-Face-Greedy) [3], [10]. Upon hitting a local maximum, the protocol switches to a different mode of operation (traversing the so-called *faces*) until it can safely revert to the greedy mode.

With another generic approach, called *Compass* [11], a routing node c forwards the packet to the neighbor x that minimizes the angle formed between c, x and the destination d. This scheme, whose performance is similar to that of *Greedy*, suffers from loops. To overcome it, a randomized variant of Compass has been suggested in [2], whereby one of two neighbors (located on different sides of the line connecting c to d) is chosen at random.

The idea of randomized selection from the two sides of the line connecting the routing node to the destination has found its way into some other schemes, notably the so-called *AB* routing family [6], [9]. Generally, those protocols differ in the assignment of weights to the next-hop candidates, which may combine measures of distance and angle at the same time. Further advancements involve more sophistications of the neighbor weights. One popular class of solutions, exemplified by PGR (Probabilistic Geographic Routing) [14], involves a randomized selection from an angular sector centered around the direction towards the destination. The weight assignment scheme of PGR includes the measured reliability of the neighbor as well as its remaining battery power.

Geographic Random Forwarding (GeRaF) [16] is another variation on the theme, whereby the regions for next-hop neighbor selection are determined by the distance of their members to the destination. A shared feature of all those schemes is that they enlarge the region size in some systematic manner upon a failure to locate an eligible next-hop neighbor within a narrower (preferred) region. In this respect, the Adaptive Load-Balanced Algorithm (ALBA) [4] prescribes a dynamic coloring scheme aimed at identifying dead ends. It also postulates modifications to the MAC scheme allowing multiple neighbors to heed the RTS packet of the forwarding node. Having received such a packet, the eligible neighbors determine their so-called *queue priority index* and respond in a way that will automatically select the winner.

Deterministic routing schemes tend to receive more attention than randomized ones, which is partly due to the focus on quality of service and delivery guarantees, incorrectly perceived as incompatible with nondeterminism [1]. To prove that a particular location–based routing scheme guarantees packet delivery is often equivalent to demonstrating that the underlying communication graph observes certain properties, in particular, that it is *planar*. The planarity can be "forced" by excluding from consideration certain (otherwise legitimate) edges/links present in the network. We note that while guaranteed delivery is a desirable property, the overhead to traffic introduced by forcibly planarizing a graph is not trivial and could be significant [12], [13]. As it happens, optimal paths in the original network graph are often eliminated from the reduced, planarized, graph. We are therefore dealing with the following paradoxical situation: guaranteed delivery protocols are non-optimal in terms of path overheads, yet their promise is to be "better" than protocols where explicit routing information needs to be calculated. This paradox is exacerbated by the fact that the postulate of reliable delivery is generally ill-defined in the ad-hoc wireless environment: it seldom means "actual" reliability, the way it has been traditionally viewed in the wired world. There are a few cases where continuous non-optimal selection (as instigated by planarized geometric routing) could be considered advantageous over a (possibly infrequent) overhead of determining routing paths (information security and path anonymity come to mind). However, for a general-purpose routing scheme, guaranteed delivery may be a lesser priority, as long as most traffic finds its way to the destination with little or no added overhead. We adopt this view, i.e., the guaranteed delivery is only one facet, and has to be seen in the context of other overheads.

III. Protocol Definitions

Formally, we define an ad-hoc wireless network as a set V of n nodes placed in two dimensional Euclidean space. Each node is aware of its location expressed as a pair of Cartesian coordinates (x, y). Two nodes can communicate with each other if and only if their Euclidean distance is at most R, which is represented by an edge between the corresponding nodes. The resulting graph, *G=(V,E)*, describes the topology of the network. Given G and a pair of nodes (i, j), $i, j \in V$, the problem of location-based routing is to find a path from i to j that maximizes packet delivery ratio while minimizing the path length without degrading the network throughput.

Suppose that the node about to make a routing decision, the next hop node, the destination, and the set of neighbors of the current node are denoted c, x, d, and $N(c)$, respectively. The Euclidean distance between two nodes a and b is denoted by $dist(a, b)$. let $\theta(i) = \angle dci$ be the angle between c, d, and one of the neighbors of c, $i \in N(c)$. Our proposed protocols can be formally described as follows:

Random Walk (RW): The candidate nodes, $FS(c) \in N(c)$, reside inside a sector of size $\pi/2$ (RW-90) or π (RW-180) centered at c and around the line cd. c forwards the packet to one $x \in FS(c)$ chosen uniformly at random. (Note that we define two variants of RW here.)

Forward-First with Rank (FFR): c defines two subsets of neighbors $NF(c)$ and $NB(c)$ such that the first set consists of neighbors from the forward sector of size π towards the destination d, while the second set contains the remaining neighbors. The next node x is chosen in the following manner;
a) if $NF(c)$ is not empty then x is the neighbor from $NF(c)$ that maximizes $r \times |\cos\theta(x)|$, where r is the distance between c and x; $\theta(x)$ is the angle formed between x, c, and d.
b) otherwise, x is the neighbor (from $NB(c)$) that minimizes $s \times |\theta(x)|$; where s is the distance between x and d.

Forward with Random selection out of Two (FRT): c chooses two candidates x_1 and x_2 from $NF(c)$ that are the first and the second best neighbors in terms of maximizing $r \times |\cos\theta(i)|$, where r is the distance between c and i; $\theta(i)$ is the angle formed between c and d. The next hop node x is then chosen uniformly at random out of x_1 and x_2. FRT-90 selects the candidates from a $\pi/2$ sector centered at c and around the line cd. For FRT-180, the sector size is π.

IV. Simulation Setup

We have evaluated our protocols under SMURPH [5]. For grid network models, we considered a 10×10 perfect grid with the edge (node-to-node link) length of 50m. The random topologies were generated by uniformly distributing 75 nodes within a 500m $\times$ 500m area. The transmission radius (i.e., the geographical range of a neighborhood) was set to 90m in all the experiments. That distance was used as a straightforward way to define neighborhoods and corresponded to the nearly maximal node separation at which the bit error rate of the channel was still acceptable. Note that our choice of the transmission radius did not affect the operation of the simulator's channel model. That measure was an internal assumption based on the *observed* properties of the wireless channel, which the nodes used to determine their neighborhoods in a purely geographical fashion.

For randomly generated topologies, only connected networks were used (occasional disconnected instances were detected beforehand and discarded). The transmission radius ensured that the vast majority of randomly generated topologies were indeed connected. Our uniform traffic model was augmented by the requirement that every packet should be subjected to non-trivial routing, i.e., the destination of every packet was at least two hops away from the source.

In the case of biased traffic, the endpoints were located on the edges (specifically the bottom and upper edge of the grid), while the interior nodes acted exclusively as routers. We considered three groups of source-destination pairs, such that every packet had to pass through the entire network width. That allowed us to have a controllable configuration of hot spots in the middle separated from each other by a few hops.

We simulated an IEEE 802.11 MAC protocol under a shadowing propagation model with the path loss at distance d equal to

$$PL(d)[dB] = PL(d_0) + 10\beta \log(d/d_0) + X_\sigma$$

where $PL(d_0)$ is the path loss at the reference distance d_0 (1m), β is the path loss exponent, and X_σ is a zero-mean Gaussian distributed random variable with standard deviation σ. In our experiments, β and σ are 3 and 1, respectively. The received power is expressed as

$$P_r[dBm] = P_t[dBm] - PL(d)[dB]$$

where P_t is the transmission power (10dBm). If the received power is less than the threshold power P_{th}, the signal cannot be received even in principle. Otherwise, the bit error rate during reception depends on the dynamic signal to interference (SIR) ratio. The physical medium bit rate was fixed to 5 Mbps (resulting from the halved Manchester rate of 10 Mbps).

V. Simulation Results

All results are the averages of 10 replications. The confidence intervals to within 95% have been determined as well but are not plotted as they are very tight. All figures are plotted in log-log scale to compress a wide range of loads and throughputs within the available space. Packets are dropped either when the routing algorithm cannot find a next hop according to the rules provides in section III, or due to buffer overflow at the node. In all simulations, the per-node buffer size was 15 packets.

A. Grid Network / Uniform Traffic

Figure 1(a) shows that RW-180 has the worst performance and Greedy has the best throughput under uniform burst traffic on grid networks. GeRaF can comfortably compete with Greedy across the entire range of loads. The shortcomings of RW-180 are obvious: the protocol selects the next node at random from all neighbors inside the sector of size π. Odds are against an optimal choice, even if by restricting the sector to π we can guarantee some "progress." On the other hand, GeRaF divides the forward sector into regions based on geographic progress towards destination; thus, chance is high that it may select the closest node. The next three protocols are RW-90, FRT-180, and FFR. Due to the limited sector size of $\pi/2$, biased randomization, and maximized projection, these three protocols might end up following similar paths. FRT-90 achieves better throughput compared to FRT-180, because of its narrower (angular) sector size, which limits the extent of deviation from the optimal path. Under non-biased conditions, this approach is always a safe bet. Thus, by the same token, RW-180 loses with respect to RW-90. Greedy, GeRaF, RW-90, FRT-180, and FFR behave roughly the same at low loads but Greedy and GeRaF prevail at high loads. These arguments also explain the average path length exhibited by the protocols and

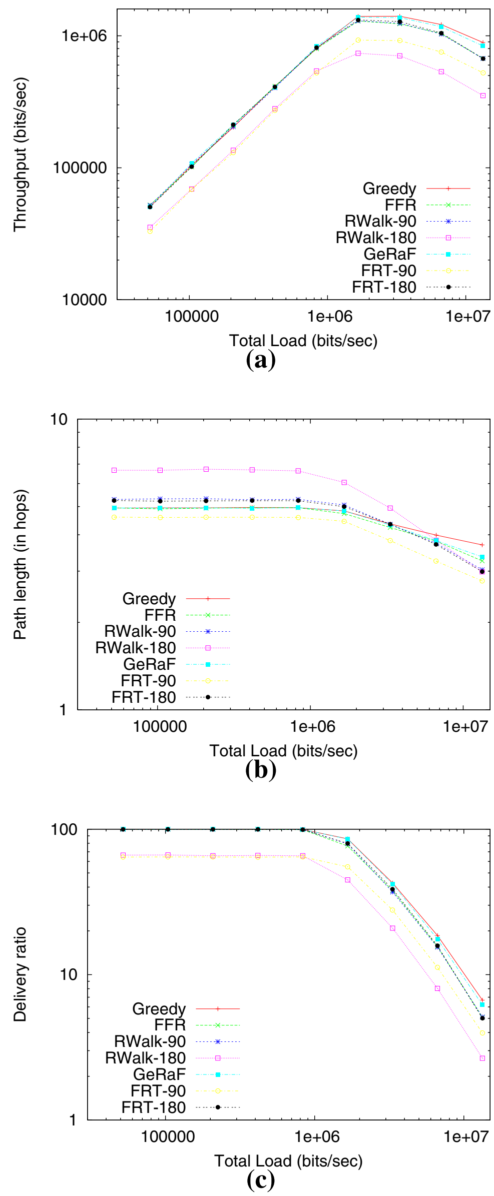

Fig. 1. (a) average network throughput, (b) path length, and (c) packet delivery ratio, for 10 $\times$ 10 grid.

shown in Figure 1(b). The packet delivery ratio (Figure 1(c)) shows a similar behavior to throughput.

B. Random Network / Uniform Traffic

Figure 2 illustrates the behavior of all protocols on random networks. Under light load, FRT-90 has the lowest throughput, whereas, greedy and GeRaF exhibit the best performance. This behavior is quite different form what we saw in the grid case. Interestingly, under heavy load, FRT-90 outperforms all other protocols, while Greedy and GeRaF yield significantly. RW-90 and FFR appear similar to FRT-90, while RW-180 turns out the loser. This can be explained by the role of the sector size. Both FRT-90 and RW-90 have a limited number of eligible neighbors, which tends to translate into short paths

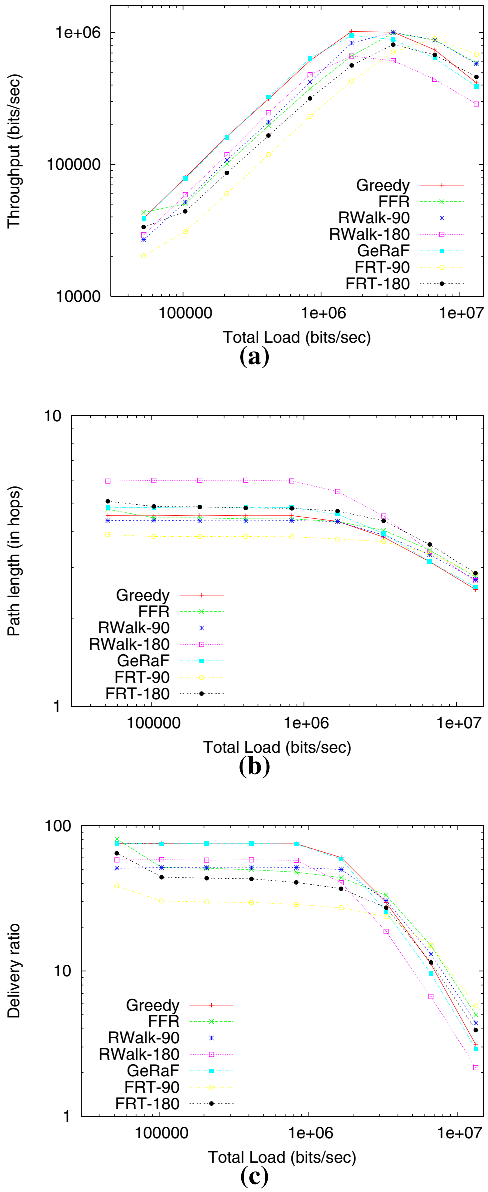

Fig. 2. (a) average network throughput, (b) path length, and (c) packet delivery ratio, for random topology.

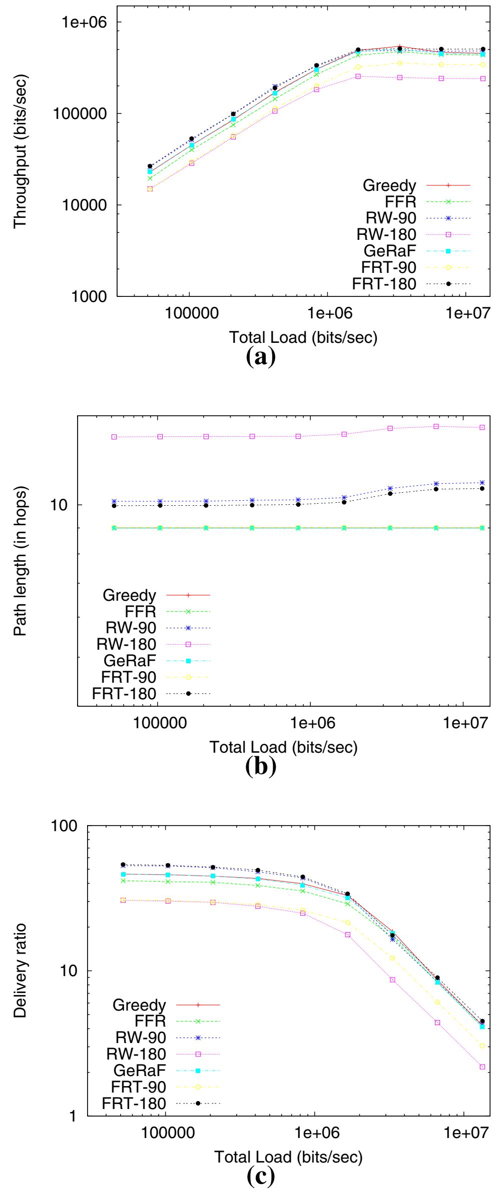

Fig. 3. (a) average network throughput, (b) path length, and (c) packet delivery ratio, for 10 × 10 grid with biased traffic.

and reasonable performance under all loads. Under light load, RW-180 performs well, because of its tendency to diversify routes and explore more paths (note that its delivery ratio is higher than for RW-90). However, it loses at heavy load, when the longer paths cause premature congestion, which in turn reduces the throughput and delivery ratio. FFR performs better than Greedy under high load, as it has the option of choosing a neighbor from the backward sector, which helps it cope with local minima.

In confrontation with Figure 1, this set of results hints at a fundamental difference between regular and actual networks. The traffic still being uniform, the accidental nature of paths in a random network provides opportunities for nondeterministic routing decisions to show their advantage over locally optimal ones. Even though the shortest path objective (minimizing the amount of network resources along the path) still seems to prevail under heavy load (where the availability of bandwidth becomes essential), we clearly see the benefits of *random exploration* at the lighter end of the traffic spectrum.

C. Grid Network / Biased Traffic

The most interesting results, from the viewpoint of their realism, have been obtained with biased traffic on a grid network, which is shown in Figure 3. The purpose of this configuration was to introduce traffic hot-spots in the middle of the grid. The odds are clearly stacked against deterministic protocols that attempt single "best" choices. Similar to the previous networks and traffic condition, we observe that the worst throughput is achieved by RW-180. FRT-90, which had some hope of competitiveness based on the two previous configurations, can evidently sustain higher saturation throughput, but its performance is similar to RW-180. FRT-90 selects the next hop in a biased random uniform fashion, and reaches the destination in fewer hops than RW-180.

The two deterministic protocols, Greedy and FFR, lack options to distribute load among neighbors and forward the packets to the congested next hop. Their ability to track the shortest path, as illustrated in Figure 3(b) is evident but at the detriment of overall throughput. Notably, FRT-90 is also exactly tracking the shortest path as a side–effect of limited next-hop options in a grid when forwarding within sectors of 90 degrees. On the other hand, the best performance is observed for two randomized protocols with notably worse hop count performance, namely RW-90 and FRT-180. The tradeoff expressed by the two protocols is quite telling of the overall strategies possible for randomized protocols. RW-90 expresses randomization without any ranks or weights that would creatively bias next node choices, but is well focused within a narrow sector. Opting for a larger sector (like RW-180) can destroy its advantages. On the other hand, FRT-180 expresses randomization with ranking of nodes based on some figure of merit that characterizes good and bad choices, but is given a larger sector to pick for its possible candidates. The two schemes achieve comparably high performance. Also, as in the uniform traffic case, GeRaF performs similar to Greedy.

Of course, nothing comes for free, and the sub-optimal choices of RW-90 and FRT-180 eventually catch up at high loads with the performance of Greedy and FFR. Their differences become small to be statistically significant, yet at ultra-high loads, FRT-180 and RW-90 hold their ground quite well: they are able to circumvent hot-spots, which Greedy and FFR fail to do. The explanation of what happens at very high loads is better provided in the context of the number of hops in Figure 3(b) where a slight increase of the number of hops is seen for RW-90 and FRT-180 at very high loads. This means that the packets that eventually make it to their destinations have followed inflated paths (possibly by only just one hop), but that allowed them to bypass the congested nodes, where they would have been dropped.

VI. Conclusions

In summary, we have observed that under biased traffic, some of the randomized schemes allow for better load balancing over a wide range of loads. Uniform randomization appears to work better if it operates within a confined sector, whereas biased randomization needs the exact opposite, i.e., larger sectors, to better distribute the load. Some of natural intuitions, like "routing randomization works well with random networks," have proved to be wrong. We have established that while randomized routing protocols can deal with random environments, one has to be aware of the additional cost that the randomization will place on path lengths and therefore congestion. The simplest location-based scheme, Greedy, is quite capable to handle a variety of scenarios except for high load non-uniform traffic situations. Randomized routing has a lot to offer for biased traffic, even in regular topologies. It remains a question whether one can develop a deterministic routing scheme that will perform well in biased traffic scenarios.

Note that, in contrast to the wired world, there is nothing to be gained in terms of throughput by applying deflection in wireless (without reusing the bandwidth of the deflected transmission). In this context, it would be interesting to analyze the protocol behavior with directional antennas, where simultaneous non-interfering communications in the vicinity of a node are possible which, one could argue, will bring the model of operation closer to that of deflection in wired networks. As part of future research, we are also interested in developing analytical models for the path length distribution of deterministic and randomized routing protocols.

References

[1] C. Baransel, W. Dobosiewicz, and P. Gburzyński. Routing in multi-hop switching networks: Gbps challenge. *IEEE Network Magazine*, (3):38–61, 1995.

[2] P. Bose and P. Morin. Online routing in triangulations. In *10th Annual International Symposium on Algorithms and Computation (ISAAC '99)*, pages 113–122, 1999.

[3] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. In *3rd Workshop on Discrete Algorithms and Methods for Mobile Computing and Communication (DIALM)*, 1999.

[4] P. Casari, M. Nati, C. Petrioli, and M. Zorzi. ALBA: An adaptive loadbalanced algorithm for geographic forwarding in wireless sensor networks. In *Military Communications Conference, 2006. MILCOM 2006*, pages 1–9, October 2006.

[5] W. Dobosiewicz and P. Gburzyński. Protocol design in SMURPH. In J. Walrand and K. Bagchi, editors, *State of the art in Performance Modeling and Simulation*, pages 255–274. Gordon and Breach, 1997.

[6] T. Fevens, I.T. Haque, and L. Narayanan. A class of randomized routing algorithms in mobile ad hoc networks. In *Algorithms for Wireless and mobile Networks (A_SWAN 2004)*, Boston, MA, August 2004.

[7] G.G. Finn. Routing and addressing problems in large metropolitan-scale internetworks. Technical Report ISU/RR-87-180, USC ISI, Marina del Ray, CA, March 1987.

[8] P. Gburzyński and J. Maitan. Deflection routing in regular MNA topologies. *Journal of High Speed Networks*, 2(2):99–131, 1993.

[9] I.T. Haque, C. Assi, and W.J. Atwood. Randomized energy aware routing in mobile ad hoc networks. In *The 8th ACM International Symposium on Modeling, Analysis and Simulation*, Montreal, October 2005.

[10] B. Karp and H. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Proc. of 6th ACM Conference on Mobile Computing and Networking (Mobicom'00)*, 2000.

[11] E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks. In *Canadian Conference on Computational Geometry (CCCG'99)*, pages 51–54, 1999.

[12] X. Li, G. Calinescu, and P. Wan. Distributed construction of planar spanner and routing for ad hoc networks. In *Proceedings of IEEE INFOCOM'02*, pages 1268–1277, New York, NY, June 2002.

[13] A. Liebers. Planarizing graphs–a survey and annotated bibliography. *Journal of Graph Algorithms Applications*, 1(5):1–74, 2001.

[14] T. Roosta, M. Menzo, and S. Sastry. Probabilistic geographic routing in ad hoc and sensor networks. In *Int. Workshop on wireless Ad-Hoc Networks*, 2005.

[15] H. Takagi and L. Kleinrock. Optimal transmission ranges for randomly distributed packet radio terminals. *IEEE Transactions on Communications*, 32(3):246–257, 1984.

[16] M. Zorzi and R. R. Rao. Geographic random forwarding (GeRaF) for ad hoc and sensor networks: multihop performance. *Mobile Computing, IEEE Transactions on*, 2(4):337–348, 2003.