On Optimal Geographic Routing in Wireless Networks with Holes and Non-Uniform Traffic

Sundar Subramanian, Sanjay Shakkottai and Piyush Gupta

***Abstract*— Geographic forwarding has been widely studied as a routing strategy for large wireless networks, mainly due to the low complexity of the routing algorithm, scalability of the routing information with network size and fast convergence times of routes. On a planar network with no holes, Gupta and Kumar (2000) have shown that a uniform traffic demand of $\Theta(1/\sqrt{n \log n})$ is achievable. However, in a network with routing holes (regions on the plane which do not have active nodes), geographic routing schemes such as GPSR or GOAFR could cause the throughput capacity to significantly drop due to concentration of traffic on the face of the holes. Similarly, geographic schemes could fail to support non-uniform traffic patterns due to spatial congestion (traffic concentration) caused by greedy "straight-line" routing.**

In this paper, we first propose a randomized geographic routing scheme that can achieve a throughput capacity of $\Theta(1/\sqrt{n})$ (within a poly-logarithmic factor) even in networks with routing holes. Thus, we show that our scheme is throughput optimal (up to a poly-logarithmic factor) while preserving the inherent advantages of geographic routing. We also show that the routing delay incurred by our scheme is within a poly-logarithmic factor of the optimal throughput-delay trade-off curve.

Next, we construct a geographic forwarding based routing scheme that can support wide variations in the traffic requirements (as much as $\Theta(1)$ rates for some nodes, while supporting $\Theta(1/\sqrt{n})$ for others). We finally show that the above two schemes can be combined to support non-uniform traffic demands in networks with holes.

I. Introduction

Geographic forwarding based techniques have been widely suggested as an efficient routing method for wireless and sensor networks [26], [11], [14]. A key advantage of geographic routing is that the nodes are not required to maintain extensive routing tables, and can make simple routing decisions based on the local geographic position of its neighboring nodes, i.e., they can choose the neighbor node that is closest to the destination and forward the packet to it. As the nodes only need to store the location of the neighbors, the routing information grows as the density of the network rather than the size of the network [12], and hence is scalable. In non-uniform networks, the geographic forwarding strategies may fail due to circumstances in which a forwarding node may not have any neighboring nodes that are closer to the destination than itself and may get stuck in routing "holes" or local minima.

S. Subramanian and S. Shakkottai are with the Wireless Networking and Communications Group, Department of Electrical and Computer Engineering, The University of Texas at Austin, {ssubrama,shakkott}@ece.utexas.edu. P. Gupta is with Bell Labs, Alcatel-Lucent, pgupta@research.bell-labs.com. This research was supported by NSF Grants CNS-0325788, CNS-0347400, CNS-0519535 and Darpa CBMANET, Darpa ITMANET programs.

While routing protocols such as [11], [4], [14] overcome the "hole" problem by switching to a boundary tracing scheme until geographic forwarding is possible, these methods typically induce a large number of packet routes to share the same spatial region around the holes, causing significant congestion along the boundaries and a consequent loss in throughput capacity. In fact, this phenomenon is common to routing algorithms that compute the shortest paths (w.r.t some metric of distance) between the source and destination nodes. Many popular MANET algorithms such as DSDV[21], AODV[22] or DSR[10] are based on geographically shortest paths or have excessive communication and packet overheads.

Alternately, routing algorithms designed for maximizing network throughput are typically dynamic algorithms involving some form of feedback and load-balancing. For example, in [27] a queue-state based packet forwarding algorithm is shown to be provably throughput-optimal. In [8], a distributed Bellman-Ford like algorithm with delay based distance metric is proposed to improve the average delay. However, a fundamental issue with load balancing based approaches is the trade-off between stability and convergence times - the algorithms may be slow to converge to good solutions, or may become unstable in the presence of delayed feedback information [2], [3]. In the rest of this paper, we restrict ourselves to static routing schemes (such as geographic forwarding) that provide fixed routes and are non-adaptive.

In the context of static routing, currently known schemes [9], [17] only allow for small variations (within $\Theta(1/\sqrt{n})$) in node data rates. However, wireless networks may demand widely varying data rates, for example, in networks with a mixture of video flows and short messaging.

In this paper, we construct a geographic forwarding based routing scheme for networks with routing holes that can support wide variations in the traffic requirements - as much as $\Theta(1)$ rates for some nodes, while supporting $\Theta(1/\sqrt{n})$ for others. To the best of our knowledge, this is the first static constructive scheme that can support such wide variations while simultaneously being throughput optimal (up to a poly-logarithmic factor).

A. Main Contributions

We consider a random planar network in which n nodes, each with circular radio range of $M(n) = \Theta(\sqrt{\frac{\log n}{n}})$, are uniformly and randomly distributed over a unit region. We allow for a finite number of constant area "holes" to occur on this network, removing any nodes that might fall within the "hole" region. We assume $\Theta(n)$ randomly chosen source-destination pairs, and define a throughput-capacity $T(n)$ as the data-rate that can be simultaneously supported between all the pairs, and the delay $D(n)$ as the mean time taken for a packet to travel from the source to its destination. Our main contributions are:

1) We study the throughput-capacity and delay performance of some geographic routing schemes in networks with holes. We show that while an upper bound on the throughput $T(n)$ is $\tilde{\Theta}(\frac{1}{\sqrt{n}})$ (see notation[1]) geographic routing schemes such as [11] can cause the capacity to drop to $O(\frac{1}{n})$.
2) We devise a geographic forwarding based random routing algorithm (RANDOMWAY) that achieves a throughput $T(n) = \tilde{\Theta}(\frac{1}{\sqrt{n}})$ (is optimal up to a poly-logarithmic factor), with a favorable delay scaling of $D(n) = \tilde{\Theta}(n)$ which lies on the optimal throughput-delay trade-off curve. We also show that the routing information in the new algorithm is scalable.
3) We consider networks with wide variations in traffic demands between source-destination pairs, where some pairs require a rate of $\Theta(1)$ while other nodes require only $\Theta(\frac{1}{\sqrt{n}})$. While currently known algorithms [17] support variations in traffic only up to $\tilde{\Theta}(\frac{1}{\sqrt{n}})$, we formulate a random routing algorithm (RANDOMSPREAD) to distribute the traffic flows uniformly over the region and show that the scheme can support *any achievable traffic demand*, up to a poly-logarithmic factor.
4) Finally, we provide a scheme to combine the two previous algorithms to support non-uniform traffic demands in networks with holes.

As our algorithms are based on geographic forwarding and are static, the convergence times are better than load-balancing based approaches such as [8], [27].

B. Related Work

Geographic routing for wireless and sensor networks has been widely studied [4], [13], [1], [11], [14] in the literature. In [26], [11], [14], [4], algorithms for routing around network holes as combinations of greedy geographic forwarding and perimeter routing or "face traversal" are presented. The fundamental idea is of planarization and face traversal when greedy routing fails due to "holes" or routing local minima. While in these schemes it is necessary to maintain the underlying planar graph structure, in [6] an efficient method to identify local minima and construct routes around the holes is provided. In [7], a two phase algorithm is proposed, in which regions where greedy forwarding is possible are identified and used for routing in the next stage.

In the context of network holes and its effects, [23] analyzes the connectivity of the network in the presence of holes and provides a condition on the topology that ensures sufficient number of edge-disjoint paths between nodes. In [12] the

[1]We define $f(n) = \tilde{\Theta}(g(n))$ if $f(n) = O(g(n)(\log n)^k)$ and $g(n) = O(f(n)(\log n)^{k_1})$ for some $k, k_1 < \infty$.

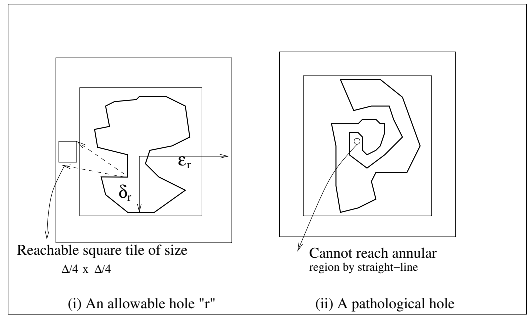

Fig. 1. Occurrence of holes in Wireless Networks

authors show that non-uniform radio patterns may induce incorrect planar graphs and can cause the routing to fail.

While the throughput capacity of networks with holes has not been explicitly studied, the results in [9] provide an upper bound on the throughput-capacity of arbitrary networks, and optimal throughput delay trade-offs are characterized in [25], [16], [5], [19]. In recent work, [17], [20] show that the throughput capacity of arbitrary networks can be studied in terms of the "min-cuts" of the network [15]. While there has been much study on efficient geographic routing methods as well as on throughput capacity of wireless networks, a systematic investigation of the effect of geographic routing strategies on network throughput and delay has not been explored previously. In this paper, we characterize the throughput-delay performance of some routing schemes discussed above and demonstrate a geography based routing algorithm that is "near-optimal" in the presence of holes and can be readily extended to non-uniform traffic requirements.

II. System Description

A. Network Model

We consider a two-dimensional model of the network in which static nodes are uniformly and randomly distributed over a unit toroidal region (to avoid edge effects). The nodes are assumed to have a uniform circular transmission range of $M(n) = \Theta(\sqrt{\frac{\log(n)}{n}})$, where n corresponds to the density of nodes in the network. Thus, $M(n)$ relates the scaling of the transmission radius to the growth in network size. The connectivity among the nodes is regulated by the transmission radius, i.e., a node is assumed to be connected to all nodes that lie within its radio range $M(n)$. It has been shown [9] that a transmission radius of $M(n) = \Theta(\sqrt{\frac{\log(n)}{n}})$ is sufficient for the network to be connected in the large-node regime, and the result assures that the number of nodes within the radio range of any node grows to *infinity* asymptotically.

To model the effect of network "holes" due to various factors such as the presence of physical obstacles, clusters of failed nodes etc., we allow for the occurrence of holes of various shapes over the unit region where the nodes are deployed. For our analysis, we consider the class of hole shapes and placements as follows. Consider Figure 1.

Assumption 2.1: Hole placements: Let δ_r be the side of the smallest square that contains the hole r, and $\epsilon_r = \delta_r + \Delta$ be the side of a larger concentric square around the hole. Then, no other hole t can be placed such that its ϵ_t outer square can intersect with that of hole r. Further, the ϵ_r outer square of any hole r cannot intersect with the boundaries of the unit square.

Assumption 2.2: Hole shapes: Consider the tiling of the unit region by square tiles of dimension $p \times p$ for some small $p > 0$. Then the holes are measurable by these tiles (they are the union of contiguous tiles). Further, any node A in the interior of the δ-square can reach any point in a square of size $\frac{\Delta}{4} \times \frac{\Delta}{4}$ in the annular region between the ϵ and the δ-squares by straight line not intersecting the hole. For an illustration, see figure 1.

Note that the fundamental problem of geographic routing with network holes (e.g. local minima) exists even in this restricted set of hole shapes. We allow for K such holes (finite number of holes, that do not scale with the network size) to be arbitrarily placed on a unit region, and assume that the nodes that fall in the interior of the holes are removed from the network. Notice that due to the restrictions on the hole placements and shapes, there is a non-vanishing fraction of the unit region that is not obscured by the holes, and hence the number of remaining nodes in the network is $\Theta(n)$ (with high probability). Also, the radio range of $M(n)$ as defined earlier is still sufficient for the connectivity of the surviving nodes (w.h.p).

B. Traffic Model

Similar to the uniform traffic model proposed in [9], we assume $n/2$ random source nodes and randomly (uniformly and independently) choose destination nodes for each traffic source node. If the source or the destination node of a traffic flow is removed due to the occurrence of a network hole, we disregard the traffic introduced into the network by such flows. We define the throughput capacity of the network as follows.

Definition 2.1: The throughput capacity $T(n)$ of a network is defined as the maximum data-rate that is simultaneously achievable by all surviving source-destination pairs.

Also, we consider the protocol model [9] to capture the interference effects of simultaneously transmitting nodes which is recalled below.

Definition 2.2: A transmission between a node A and its receiving node B is assumed to be successful if $d(A,B) \leq M(n)$ and $d(C,B) > (1+\gamma)M(n),$ for some $\gamma > 0$, for all other transmitting nodes $C \neq A$.

We define the packet delay $D(n)$ as the average time taken by the routing algorithm to travel from the source to its destination averaged over all source-destination pairs. Since packets can travel only distances lesser than the radio range in any single step, communication between any source-destination pair is through multi-hop packet relaying. Thus, the average delay for a packet can be seen as the MAC delay at each hop summed over all hops in the packet-route. When the traffic patterns are uniform over the network, the queuing delay at intermediate hops is uniform for all flows (and hops) and the packet delays are proportional to the number of hops.

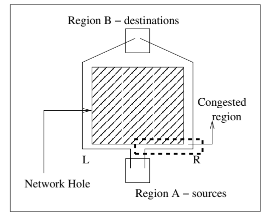

Fig. 2. Congestion around the boundaries - Effect of perimeter routing.

III. Loss of Throughput with Perimeter Routing

In this section, we study the throughput-capacity properties of some location-based routing schemes. Many geographic location based routing schemes such as GPSR [11], GOAFR[14], GEDIR[26] utilize perimeter or face routing based strategies to route around network holes. The representative idea behind these routing strategies is described below.

(i) Packets containing the position of the destination nodes are forwarded greedily to neighboring nodes that are closer to the destination.
(ii) When greedy geographic forwarding fails due to nodes that do not have any neighbor nodes closer to the destination than itself (the node is a local minima), the routing schemes switch to a perimeter-routing mode.
(iii) In this mode, a node A on receiving a packet from another node B, checks to see if it is closer to the destination than B. If yes, it reverts back to a greedy forwarding scheme. Else, it sweeps counter clock-wise from the direction $\overrightarrow{AB}$ and identifies neighbor-node C as the first node found in this search. It then chooses C as its next-hop neighbor.

Fundamentally, the basic strategy common to many such routing strategies is to follow the boundary of the hole until greedy forwarding is possible. While these strategies are scalable with respect to routing information (nodes only need to store location information of the neighboring nodes), they cause significant amounts of network congestion along the boundaries of the network holes, since the routing scheme requires that all flows with the source and destination across the network hole be routed around the boundary.We formally show that even with only one simple shaped hole in the network, GPSR based (face-routing) strategies cause a significant drop in throughput capacity.

Theorem 3.1: Consider a single square hole (as in figure 2) at the center of the unit region, with finite area. Then, under the protocol model and uniform traffic assumption, the throughput $T(n)$ that can be supported for GPSR-like strategies is $T(n) = O(\frac{1}{n})\,(a.s)$. Further, the average delay $D(n) = \Omega(n^{3/2}W(n)) \quad (a.s)$, where $W(n)$ is the size of the packet scaling with the network size n.

Proof: Due to space constraints, we only provide a sketch of the proof. We show that a sizeable fraction of the traffic have sources and destination nodes are on the opposite sides of the holes and demonstrate that GPSR-like routing strategies induce all packets to flow through the region in the vicinity of the boundary of a network-hole, causing a reduction in the throughput-capacity. Consider the subset of source-destination pairs with source nodes in region A and their corresponding

Field Name	Functionality
WAYPOINT-NUM	Number of waypoints to traverse before reaching destination
NEXT-DEST	Location of the next waypoint
FINAL-DEST	Location of the original destination
DATA	Message to the destination node

TABLE I

FIELDS IN THE HEADER OF THE PACKET.

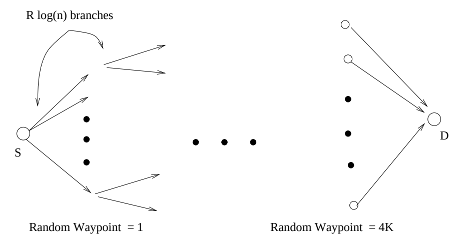

Fig. 3. The branching structure of the packet before reaching the destination.

destinations in region B (see figure 2). Since the regions A and B have a non-vanishing fraction of area, the number of such source-destination pairs is $\Theta(n)$ (with high probability) as the source and destination nodes are uniformly distributed over the unit region.

From the construction of the regions and the network hole, it follows that none of the traffic flows have greedy geographic paths to their destinations. Notice that in all schemes that utilize perimeter-routing, all the traffic flows travel through the narrow region (with a thickness of $M(n)$, the radio range) around the edge of the boundary in a counter-clockwise direction. As $\Theta(n)$ flows have to travel through the boundary, assume WLOG that the $\Theta(n)$ routes pass through the narrow strip of length $\Theta(1)$ on the right (as indicated in figure 2), and consider any tile of size $M(n) \times M(n)$ on this strip. As the protocol model allows only one packet within the tile to transmit in any given time-slot, the best achievable throughput capacity is $\Theta(\frac{1}{n})$. Further, for non-vanishing fraction of the traffic through the strip, the number of hops for any packet through this crowded strip is $\Theta(\frac{1}{M(n)})$ and the delay at each hop is $\Theta(nW(n))$, where $W(n)$ is the packet scaling. It follows that the average delay $D(n) = \Theta(n\frac{1}{M(n)}W(n)) = \Omega(n^{3/2}W(n))$. Thus, the delay due to GPSR like strategies is not on the optimal throughput-delay curve (by setting $W(n) = T(n)$ we can compare with results in [5]). ■

Remark 3.1: Note that the above result can be generalized to any hole that contains a square region of non-zero area (this includes “allowable” holes in Section II).

IV. RANDOMWAY(n, K) ALGORITHM

In this section, we describe our randomized multipath routing algorithm that can achieve near-optimal throughput-capacity, even in the presence of network holes. The algorithm takes as input the number of nodes in the network, the packet to be sent, as well as the number of holes.

The packet, in addition to the data payload and the destination location, is assumed to have a few extra fields for facilitating our algorithm. These fields are provided in Table IV. Notice that the size of the packet does not grow with the size of the network. Consider the first packet in all the source nodes. The algorithm is as follows:

1) The source node for every traffic flow creates $R\log(n)$ copies of its packet to send. It chooses $R\log(n)$ independent and uniformly distributed points from the unit region and sets the NEXT-DEST field to the randomly generated location in each of these copies. The WAYPOINT-NUM is set to $4K + 1$ in all the packet copies.
2) The $R\log(n)$ packets are routed from the source in a greedy geographic manner to the location in NEXT-DEST.
3) A node, on receiving a packet, checks if it is the NEXT-DEST location. If it is not the NEXT-DEST location, (i) it searches within its neighboring nodes for the node that is closest to the NEXT-DEST location, and forwards the packet to that node. (ii) If none of its neighbor nodes are closer to the NEXT-DEST than itself, the node drops the packet.
4) If it is the NEXT-DEST location, (i) it checks if WAYPOINT-NUM > 1. If yes, it sets WAYPOINT-NUM = WAYPOINT-NUM - 1, and makes $R\log(n)$ copies of the packet and again generates uniform and randomly chosen locations for the NEXT-DEST in each of the packet copies, and forwards them greedily. (ii) If WAYPOINT-NUM = 1, the node sets NEXT-DEST = FINAL-DEST, WAYPOINT-NUM = 0 and forwards the packet greedily. (iii) If WAYPOINT-NUM = 0, the packet is received at the destination.

Thus, the algorithm creates $R\log(n)$ copies of the first packet at the source and sends each of them to a random waypoint by greedy geographic routing. If the greedy forwarding fails due to a network hole, the packet is dropped. The packet on reaching the random waypoint node, creates $R\log(n)$ further copies and sends each of them to their randomly chosen waypoints. Thus, we create a branching tree of random waypoints, of depth $4K + 1$ and degree $R\log(n)$ (see figure 3). Note that each copy of the packet travels greedily to $4K$ intermediate destinations before it reaches its original destination. Subsequent packets follow the same route as the first packet.

V. ANALYSIS OF RANDOMWAY ALGORITHM

In this section, we show that our algorithm achieves a throughput capacity that is only a logarithmic factor away from the best-case capacity for a network with holes. We also show that our algorithm provides bounded delay that is comparable to the delay incurred in a network without holes and with straight-line routing, i.e., it is order-wise delay optimal. Further, we show that the routing information that needs to be stored in the nodes does not increase appreciably with the network size, i.e, the routing information remains scalable.

A. Throughput Optimality

In order to compare the throughput-capacity performance of our routing algorithm, we first provide a general upper bound on the the best-case capacity in networks with holes, and then show that throughput achieved by our scheme is only smaller by a poly-logarithmic factor.

Theorem 5.1: Consider a uniform random planar network, with K allowable holes in it and assume a uniform traffic pattern (as described in Section II). Then, under the protocol model for interference, the best case throughput-capacity of the network $T(n)$ satisfies $T(n) = O(\frac{1}{\sqrt{n}})$.

The above theorem is a restatement of the result in Thm 2.1 [9] where it is shown that an upper bound on the transport capacity of any arbitrary network is $\Theta(\sqrt{n})$ bit-meters per second. Since by our uniform traffic model, the source and destinations are a non-vanishing distance away from each other, it follows that by distributing this transport capacity to the $\Theta(n)$ flows, the data-rate that can be simultaneously achieved by all the flows can be no more than $\Theta(\frac{1}{\sqrt{n}})$.

We now demonstrate that a throughput-capacity of $\Theta(\frac{1}{\sqrt{n}(\log(n))^P})$ for some $P < \infty$ is *achievable* by our algorithm. To avoid technical complications due to edge effects, we assume that the network is a unit toroidal region. Note that with this assumption, the network nodes/tiles are symmetrically distributed with respect to the traffic patterns.

Theorem 5.2: Let G be a random network over a unit torus with K (a finite number) *allowable* holes placed arbitrarily, and uniformly distributed traffic flows. Then, the randomized algorithm RANDOMWAY(n, K) achieves (almost surely) a throughput capacity of $T(n) = \tilde{\Theta}(\frac{1}{\sqrt{n}})$ simultaneously for all source-transmitter pairs, under the protocol model for interference.

Proof: The proof follows in three steps. (i) We show that for all sources S and their corresponding destinations D, the algorithm ensures that at least one packet-route from S reaches the destination D via the $4K$ intermediate destinations. (ii) We construct a tiling of the unit region with tiles of side $M(n)$ and show that the number of packet routes through any such tile is upper bounded by $\tilde{\Theta}(\sqrt{n})$. (iii) We demonstrate a scheduling scheme that can achieve a throughput of $\tilde{\Theta}(\frac{1}{\sqrt{n}})$ for each surviving packet route.

Proof of (i): Consider a tiling of the remaining area of the unit region (after the placement of the holes) by tiles of size $\frac{\Delta}{4} \times \frac{\Delta}{4}$. From Assumption 2.1 on the placement of holes, we see that the tiled regions will remain connected in the presence of holes. That is, there exists a sequence of contiguous tiles to travel from any tile to any other. Since 3 straight-line paths are sufficient to go around any allowable hole (see figure 1), we have the following claim.

Claim 1: Given any source S and destination D, there exist tiles $T_0, ..., T_{4K+1}$ such that S lies in T_0, D lies in T_{4K+1} and tile T_i is *reachable* (i.e., a straight-line path that does not intersect a hole exists) from T_{i-1} for all $i \in \{1, \cdots, 4K+1\}$. (Even if there are less than K holes in between, it is possible to split a straight-line path into smaller straight-line paths so that there are exactly $4K + 1$ tiles between the two nodes.)

From our assumption on the hole shapes (Assumption 2.2) there is a $\frac{\Delta}{4} \times \frac{\Delta}{4}$ tile such that D can be reached from any point within this tile. Similarly, there is a tile such that any point within this tile can be reached from S. Without loss of generality, we assume that the tiles are T_1 and T_{4K} respectively.

From step 4 of our algorithm where a random waypoint at depth $4K + 1$ depth greedily forwards the packet to the final destination, it follows that the probability that no path is created to the destination D, $\mathbb{P}\big(\text{No path to } D\big) \leq \mathbb{P}\big(A_{4K+1}\big)$ where A_L is the event {no surviving waypoints of depth L in T_L} (i.e, those not killed in Step 3(ii)). Notice that

$$\begin{aligned}\mathbb{P}(A_{4K+1}) = \quad & \mathbb{P}(A_{4K+1}|A_{4K})\mathbb{P}(A_{4K}) + \\ & \mathbb{P}(A_{4K+1}|A_{4K}^c)(1 - \mathbb{P}(A_{4K})) \quad (1)\\ \leq \quad & \mathbb{P}(A_{4K}) + \epsilon_{4K}(1 - \mathbb{P}(A_{4K})) \quad (2)\end{aligned}$$

where ϵ_{4K} is an upper bound on the probability that no random waypoint of depth $4K+1$ was chosen in T_{4K+1} given that there was one in T_{4K}. Since we choose $R\log(n)$ points independently at random, the probability that T_{4K+1} was not chosen, i.e.,

$$\mathbb{P}(A_{4K+1}|A_{4K}^c) \leq \big(1 - (\frac{\Delta^2}{16})\big)^{R\log(n)} \leq \frac{c_1}{n^R}, \quad (3)$$

for some $c_1 > 0$. Thus,

$$\mathbb{P}(A_{4K+1}) \leq \mathbb{P}(A_{4K}) + \frac{c_1}{n^R}. \quad (4)$$

Note that as the bound on $\mathbb{P}(A_L|A_{L-1}^c)$ (similar to (3)) is independent of L, we can recursively use Equation 4 to show that $\mathbb{P}(A_{4K+1}) \leq \frac{c_2}{n^R}$, for some $c_2 > 0$. By a union bound over all the source destination pairs, and for $R > 4$ we see that

$$\mathbb{P}(\cup_{i=1}^n \{\text{No path between } S_i, D_i\}) \leq \frac{c_2}{n^2} \quad (5)$$

and hence by Borel-Cantelli's lemma (i) is almost surely true.

Proof of (ii): We construct a tiling of the unit region by tiles of size $M(n) \times M(n)$. Consider the scenario where all the nodes removed by the hole placements are reintroduced in the network, i.e, they are allowed to have their own traffic and also forward packets from other sources. Then, given any tile, the RANDOMWAY algorithm would only create more "lines" (or packet routes) than the scenario when the nodes were removed by the holes. This occurs because (i) the number of tiles covered by a source's packet is only increased by removing the holes as RANDOMWAY algorithm drops packets on hitting a hole, (ii) the reintroduced nodes offer additional traffic that increase the number of packets. We show that even in this scenario, the number of paths that pass through any tile is bounded above by $\tilde{\Theta}(\sqrt{n})$.

Let $X_L^i(S_i, D_i)$ be the i^{th} random waypoint at depth L created between the source S_i and destination D_i by RANDOMWAY(n,K). Let $\overline{(A, B)}$ be the line segment joining the points A and B. We define $\mathcal{C}(S_i, D_i)$ as the set of all line segments created by our algorithm for routing packets between S_i and D_i. That is, $\mathcal{C}(S_i, D_i) = \big\{\overline{(X_{L-1}^i, X_L^j)}$ $\forall j \in \{R\log(n) * (i - 1), \cdots, i * R\log(n) - 1\}$, $\forall i \in \{1, \cdots, (R\log(n))^{L-1}\}, \forall L \in \{1, \cdots, 4K+1\}\big\}$.

Let $G_j(i)$ be a Bernoulli random variable with $G_j(i) = 1$ if tile j, $j \in \{1, \cdots, \frac{1}{M(n)^2}\}$ was touched by any line element of $\mathcal{C}(S_i, D_i)$ $i \in \{1, \cdots, n\}$ [2]. By symmetry of the uniform traffic pattern assumption over the unit torus, all tiles are equally likely to have been touched by $\mathcal{C}(S_i, D_i)$. We now construct a collection of i.i.d Bernoulli random variables $\tilde{G}_j(i)$, $j \in \{1, \cdots, \frac{1}{M(n)^2}\}$, $i \in \{1, \cdots, n\}$ with

$$\tilde{G}_j(i) = \begin{cases} 1 & \text{w.p } \alpha(n) \\ 0 & \text{w.p } 1 - \alpha(n) \end{cases}$$

where $\alpha(n)$ is chosen to satisfy

$$\alpha(n) \quad \geq \quad \frac{\text{Total tiles touched by any line in } \mathcal{C}(S_i, D_i)}{\text{Total number of tiles}} \tag{6}$$

Since the $G_j(i)$ and $\tilde{G}_j(i)$ are Bernoulli random variables, and $P(G_j(i) = 1)$ is less than $P(\tilde{G}_j(i) = 1)$ (by construction, and the definition of $\alpha(n)$), we have that $G_j(i) \leq_{st} \tilde{G}_j(i)$, for all j, i, where $\leq_{st}$ denotes stochastic ordering [24].

Observing that the total number of lines in $\mathcal{C}(S_i, D_i)$ is $(R\log(n))^{4K+1}$ and that no line can cover more than $2\sqrt{n}$ tiles, Equation 6 is satisfied by choosing $\alpha(n) = \frac{2(R\log(n))^{4K+1}}{\sqrt{n}}$.

Using the above construction, we show an upper bound on the the number of paths passing through any tile. Given any tile j, let $H(j)$ be the number of source-destination pairs that generate a line that touches tile j. Note that for any given constant $\lambda(n)$, $P\Big(H(j) > \lambda(n)\Big) = P\Big(\sum_{i=1}^{n} G_j(i) > \lambda(n)\Big)$ and $\leq^{(b)} P\Big(\sum_{i=1}^{n} \tilde{G}_j(i) > \lambda(n)\Big)\Big)$. Notice that for any given j, $(\sum_{i=1}^{n} \tilde{G}_j(i))$ is a sum of independent random variables, and that for each i, j, $\tilde{G}_j(i)$ stochastically dominates $G_j(i)$. From Theorem 1.A.3 of [24], $\sum_{i=1}^{n} \tilde{G}_j(i)$ stochastically dominates $\sum_{i=1}^{n} G_j(i)$, and inequality (b) follows.

By the bound on sums of i.i.d Bernoulli random variables [18], $P\Big(\sum_{i=1}^{n} \tilde{G}_j(i) > (1+\beta)nE(\tilde{G}_j(1))\Big) \leq e^{-\beta^2 \, n \, E(\tilde{G}_j(1))/2}$.

By our definition of $\tilde{G}_j(i)$ in Equation 6, $E(\tilde{G}_j(1)) = \frac{2(R\log(n))^{4K+1}}{\sqrt{n}}$, and by choosing $\beta = 2$, we get

$$\begin{aligned} P\Big(\sum_{i=1}^{n} \tilde{G}_j(i) > 2\, n\, E(\tilde{G}_j(1))\Big) &\leq e^{-n\, E(\tilde{G}_j(1))/2} \\ &\leq e^{-(R\log(n))^{4K+1}\sqrt{n}}. \end{aligned} \tag{7}$$

Thus, the probability that $H(j)$ was greater than $2(R\log(n))^{4K+1}\sqrt{n}$ is exponentially small. Since the total number of lines created between any source destination pair is $|\mathcal{C}(S_i, D_i)| = (R\log(n))^{4K+1}$, the number of paths passing though any tile j is at most $H(j) * (R\log(n))^{4K+1} = \sqrt{n}(R\log(n))^{8K+2}$.

Scheduling scheme (iii): Consider a time interval of length $\mathcal{T}$. By our protocol model, a transmitting node in a tile prevents only at most a fixed number J of neighbors from transmitting simultaneously. From the technique used in [9], each tile can be colored with one of $J+1$ colors such that no two interfering tiles have the same color. Thus, each tile can transmit for a fixed fraction $\mathcal{T}/(J+1)$ of the interval. Since the number of packet routes is no more than $\sqrt{n}(R\log(n))^{8K+2}$, each route can be provided a fraction $\frac{\mathcal{T}}{(J+1)\sqrt{n}(R\log(n))^{8K+2}}$ of the time and hence a throughput of $T(n) = \Theta(\frac{1}{\sqrt{n}(\log n)^P}) \quad P < \infty$ is achievable. ■

B. Delay Properties

In recent research, [5], [16], [19] have characterized the best-achievable capacity-delay trade-offs for static wireless networks. It is shown that in networks without mobility, the best achievable throughput-delay tradeoff is $D(n) = \Theta(nT(n))$, when packet sizes scales proportionally with the throughput. We note that for network with holes, the above relation provides an upper bound on the optimal throughput-delay trade-off, as routing is restricted to the class of algorithms that do not allow packets to travel through the “hole” regions.

Here, we show that a delay $D(n) = \Theta(n(\log(n)^P W(n))$ for some $P < \infty$ is achievable with our algorithm for packets of size $W(n)$, which is only a logarithmic factor greater than the optimal delay achievable when packet sizes are scaled as in [5].

Theorem 5.3: The average packet delay $D(n)$ for a packet of size $W(n)$ between any source S and destination D is upper bounded by $\Theta(n\, W(n)(\log(n)^P)$, for some $P < \infty$.

Proof: By (i) of Theorem 5.2, there exists a path between every source and destination. Since any path is a concatenation of $4K+1$ lines, the number of tiles traveled by a packet is no more than $\frac{(4K+1)\sqrt{2}}{M(n)}$. Also, by (ii) of Theorem 5.2 the number of paths through any tile is no greater than $\sqrt{n}(\log n)^P$. Since each packet needs to wait only for $(J+1) * \sqrt{n}(\log n)^P\, W(n)$ in each tile and number of tiles to travel is at most $\frac{8K+2}{M(n)}$, it follows that the packet is delayed by no more than $\Theta(n\, W(n)(\log(n)^P)$ seconds. ■

This result shows that the delay performance of our RANDOMWAY algorithm is away from optimal only by a polylogarithmic factor.

C. Scalability of routing information

In routing schemes that operate with greedy forwarding alone (including boundary tracing/perimeter routing), the amount of routing information that is required at a node is only the location of its neighboring nodes. In our network model, the number of neighbor nodes for any nodes is atmost $\Theta(\log n)$ as the radio range is $\Theta(\sqrt{\frac{\log n}{n}})$. For the RANDOMWAY algorithm, the requirement of routing information is increased, as the nodes that are way-points for any packet need to remember the corresponding next random way-points ($R\log n$ of them) for that packet route. However, we see that this increase is not significant.

Notice that the total number (over all packet routes) of random waypoints are $n \times (R\log n)^{4K+1}$ and are uniformly distributed over the unit region. This implies that there are no more than $\Theta((\log n)^{4K+1})$ way-points in any tile of side $M(n)$. In the worst-case scenario, all the way-points might

[2] Although the random variable G is a function of the network size, we do not explicitly denote this, for notational ease.

be chosen to be at the same node, in which case the routing information it needs to store is $\Theta(\log n) + \Theta((\log(n))^{4K+1} \times R \log n)$ where the first term is the routing information to store the neighbor locations, and the second term the next waypoints for all the packets that chose the node as a way-point. Thus, the routing information for any node is $\Theta((\log n)^{4K+2})$.

VI. Routing for Non-uniform Traffic Patterns

While the analysis in the previous sections had assumed a uniform traffic pattern, in many scenarios the traffic demands could be non-uniform and the requirements may vary widely from one node to another. For example, such patterns could be seen in a large network where there could be flows demanding much larger bandwidth than others (e.g., a mixture of video flows and short messaging). In recent research, approaches have been made to characterize the traffic patterns that can be supported in a random planar network. In [17], authors demonstrate that variations in traffic demand of the order of $O(\frac{1}{\sqrt{n}})$ are supportable, by using a Valiant-Brebner [28] scheme to distribute a source's load to all other nodes, and then solving a uniform multicommodity flow (UMF) problem [15]. We also note that the solution to this problem is nonconstructive and is based on the dual graph of the network. We however note that the motivation in [18] is different from ours. The objective in [18] is to study *a wide class of network models* by reducing them to UMF problems. However, this approach only supports "small variations" in traffic rates.

On the other hand, we provide a constructive scheme (RANDOMSPREAD) to distribute the traffic flows uniformly over the region and show that the scheme can support *any* achievable traffic demand (including $\Theta(1)$ variations in traffic), up to a poly-logarithmic factor. To the best of our knowledge, this is the first constructive scheme that can support such wide variations while simultaneously being throughput optimal (up to a poly-logarithmic factor).

Let Λ_{ij} be the traffic demand between source i and destination j and the traffic matrix Λ define the traffic demands of the network. Then, we show that if a traffic requirement Λ is feasible under any routing algorithm, the RANDOMSPREAD algorithm can achieve the rate matrix Λ (up to a poly-logarithmic factor less). Also, the advantageous properties of the original algorithm viz. bounded delay, minimal per-node routing information and robustness to location errors are preserved. For simplicity in presentation, we use the following simplified two level traffic model to analyze the performance and do not consider the presence of holes in the network. Later, we show that this can be extended to a general traffic model.

Assumption 6.1: Traffic Model: We consider a network on the unit torus, with n nodes uniformly and randomly distributed over it. There are no holes in this network. The traffic is generated by $\frac{n}{2}$ randomly chosen source-destination pairs with the following properties.

1) Each source-destination pair is generated by throwing a line randomly on the unit region, the length of the line could come from any arbitrary distribution on $[3\epsilon, 1]$.

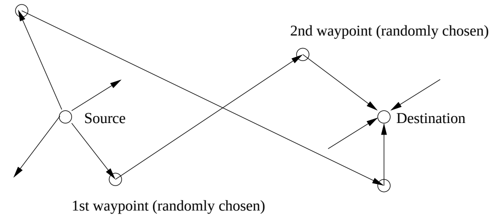

Fig. 4. Two way-point routing for non-uniform traffic

2) Each source could either be a Type-a source with a traffic requirement of $\Theta(\frac{1}{\sqrt{n}})$ or Type-b source with a requirement of $\Theta(1)$.
3) The distribution of the two kinds of loads over the source nodes is *arbitrary*.
4) The distribution is such that there exists a feasible routing scheme that supports the traffic patterns.

RANDOMSPREAD(n,ϵ): Type-a nodes send their packets directly to the destination by greedy geographic forwarding. Due to the absence of holes in this model, greedy forwarding is successful. Type-b nodes create $\sqrt{n}$ routes simultaneously to the destination, each using a three meta-hop path, as shown in figure 4. That is, for the first packet in each route, the source chooses a 1st and a 2nd waypoint by throwing a line at random and the packet is finally routed back to the destination from the 2nd waypoint. Subsequent packets follow the route of the first packet. We now show that the algorithm is optimal.

Theorem 6.1: Let Λ be a traffic matrix satisfying the properties of our traffic model. Then, algorithm RANDOMSPREAD achieves a rate of $\tilde{\Theta}(\Lambda_{i,j})$ (a.s), for all source-destination pairs (i, j).

Proof: The proof technique is as follows. We tile the unit torus region by tiles of side $M(n)/4$. We show that the number of packet routes through any tile is no more than $\tilde{\Theta}(\sqrt{n})$. By the scheduling scheme (part (iii) in proof of Theorem 5.2) this allows a throughput of $\tilde{\Theta}(\frac{1}{\sqrt{n}})$ per each packet route. Since the RANDOMSPREAD algorithm increases the number of packets simultaneously transmitted by type-b sources (by a $\sqrt{n}$ factor), the throughput achieved by the type-b sources is $\tilde{\Theta}(\sqrt{n}\,\frac{1}{\sqrt{n}}) = \tilde{\Theta}(1)$. It remains to be shown that the number of packet routes through any tile is indeed upper bounded by $\tilde{\Theta}(\sqrt{n})$.

We now partition the packet-routes in the network into 4 disjoint classes.

T_1: Packet routes generated by type-a source nodes to their corresponding destinations.

T_2: Outward lines radiating from type-b source nodes to their first intermediate way-point.

T_3: Inward lines radiating into type-b destination nodes from their last intermediate way-points.

T_4: The rest of the packet routes generated between the first and the last intermediate way-points for type-b source's packets.

The number of routes through any tile is the sum of the routes of each class T_i, $1 \le i \le 4$, through it.

Claim 2: T_1 and T_4 are $\Theta(\sqrt{n}(\log(n))^P)$, for some $P < \infty$ (almost surely).

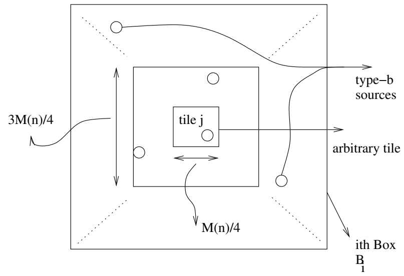

Fig. 5. Concentric tiles in the proof of Theorem 6.1

Proof: For T_1: Note that as the requirement of type-a nodes is $\Theta(\frac{1}{\sqrt{n}})$, only one packet (line) per node is sent. Since the number of type-a nodes is less than $n/2$, and the source and destination nodes are randomly distributed, the load per tile is obtained by dropping $n/2$ lines (according to the traffic model) randomly on a unit torus. The bound on the number of lines through a tile of $M(n) \times M(n)$, by dropping $\Theta(n)$ lines is at most $\Theta(\sqrt{n \log(n)})$ (from Lemma 4.13 of [9]).

For T_4: By the upper bound on the transport capacity [9], $\Theta(\sqrt{n})$, it follows that there can be at most $\Theta(\sqrt{n})$ type-b nodes. Each of these nodes send out $\sqrt{n}$ packets and choose a random second and third waypoint (see figure 4) for each packet. These waypoints are chosen by dropping a random line (with property 1 of Traffic model) and selecting the end points. Thus, the T_4 traffic generated by type-b nodes is equivalent to throwing n lines randomly. Again, by Theorem of [9] it follows that the T_4 traffic per tile is no more than $\Theta(\sqrt{n \log n})$. ■

Noticing that T_3 is analogous to T_2, it is now sufficient to show that the traffic T_2 in any tile is $\tilde{\Theta}(\sqrt{n})$. Consider any tile j and construct concentric squares around it as in figure 5. We define the following collection of sets as follows. let $\mathcal{B}_1 =$ tile j.

$$\mathcal{B}_i = \quad \{\text{Neighbors of all tiles in } \mathcal{B}_{i-1}\} \cup \mathcal{B}_{i-1}. \qquad (8)$$

Let l_i be the number of type-b sources inside $\mathcal{B}_i$.

Claim 3: $l_i \le c_1(2i-1), \quad \forall i = 1, 2, \cdots, \frac{\epsilon}{M(n)}$, for some $c_1 > 0$.

Proof: From Assumption 6.1 (property (1)), the destination node for each type-b source within any box $\mathcal{B}_i$ is outside it. Also, from (4) of Assumption 6.1, since the traffic distribution is achievable, the traffic demand of the nodes inside $\mathcal{B}_i$ cannot exceed the min-cut capacity of edges leaving the $\mathcal{B}_i$. Since each tile can at most support a constant throughput of $c_2 > 0$ and the number of tiles in the perimeter is no more than $4\sqrt{n} \times (2i-1)$ (the tiles on the boundary of $\mathcal{B}_i$ are the tiles that can transmit across the perimeter). Since each type-b node has a traffic demand of $\sqrt{n}$, the number of such nodes inside $\mathcal{B}_i$ cannot exceed $c_1(2i-1)$ for some $c_1 > 0$. ■

Now we consider the T_2 load due to l_i, $1 \le i \le \frac{\epsilon}{M(n)}$ on tile j, i.e., we count the number of lines (packet-routes) through j due to all the type-b nodes within $\mathcal{B}_{\frac{\epsilon}{M(n)}}$. (We shall show later that the load due to all type-b nodes outside the $\mathcal{B}_{\frac{\epsilon}{M(n)}}$ is also of the same order).

Recall that there are n nodes uniformly distributed in the network. We construct sectors of angular separation $\frac{2\pi}{\sqrt{n}}$ around each node with a common $0°$ angle (x-axis) for all nodes (i.e., there are $\sqrt{n}$ sectors for each node). Consider any node A, and suppose $\sqrt{n}$ random lines radiate outwards (i.e., the destination end of each of the lines is uniformly random). We observe that the probability that there are more than $W \log n$ lines in any one of the sectors is exceedingly small ($\sim \frac{1}{n^W}$). Since there are only a total of n nodes (and thus, at most n type-b nodes), and each node has $\sqrt{n}$ sectors, it now follows (using an union bound) that the number of lines radiating outwards from any of the (type-b) nodes and through any of its corresponding sectors is uniformly bounded by $W \log n$, with a probability that decays at least as fast as $1/n^{W-1.5}$. Choosing $W = 5$, and from Borel-Cantelli Lemma, we have the above property holding almost surely. Thus, without loss of generality, in the rest of the proof, we will assume that the maximum number of lines radiating outwards from any type-b node and through any of its corresponding sectors defined above (each of angle $\frac{2\pi}{\sqrt{n}}$) is no more than $5\log(n)$.

Now, consider a type-b node at a distance i from the tile j (i.e, a node in $\mathcal{B}_i$ but not in $\mathcal{B}_{i-1}$) and let ρ_i be the number of lines through tile j due to this type-b node (i.e, the number of outward radiating lines from this type-b node that intersects tile j). Then, from the discussion above and straightforward geometric arguments (essentially, counting the number of sectors of the type-b node at distance i that can "cover" tile j),

$$\rho_i \le \left(\frac{\alpha \log n}{2i-1}\right)\sqrt{n}, \qquad (9)$$

for all $1 \le i \le \frac{\epsilon}{M(n)}$, for some (finite) fixed $\alpha > 0$.

Let a_i be the number of type-b sources in $\mathcal{B}_i - \mathcal{B}_{i-1}$. Then, $l_i = \sum_{k=1}^{i} a_k$. The total number of lines through tile j due to type-b source nodes within box $\mathcal{B}_{\frac{\epsilon}{M(n)}}$ is lesser than

$$\sum_{k=1}^{\frac{\epsilon}{M(n)}} \rho_k a_k, \quad \text{s.t.} \sum_{k=1}^{i} a_k \le c_1(2i-1). \qquad (10)$$

We provide an upper bound on the number of lines through a tile by maximizing the sum in Equation 10 as follows.

$$\max_{a_i, 1 \le i \le \frac{\epsilon}{M(n)}} \alpha\sqrt{n}\log n \sum_{i=1}^{\frac{\epsilon}{M(n)}} \left(\frac{a_i}{2i-1}\right) \text{ s.t } \sum_{k=1}^{i} a_i \le c_1(2i-1)\, \forall i. \qquad (11)$$

Using standard optimization techniques, we can bound the above ILP (by an LP relaxation) to obtain the following result (we skip the proof due to space constraints).

Claim 4: The solution $a^* = c_1(1, 2, 2, 2, \cdots, 2)$ provides an upper bound for the cost function in Equation 11.

Substituting this optimal solution in Equation 11 and observing that the cost function grows as the sum of a harmonic series (and hence $\Theta(\log n)$), the number of lines through tile j due to sources within $\mathcal{B}_{\frac{\epsilon}{M(n)}}$ can be no more than $\tilde{\Theta}(\sqrt{n})$ (a.s).

We now consider the effect of all type-b sources outside the $\mathcal{B}_{\frac{\epsilon}{M(n)}}$ box. Let $\tilde{a}$ be the number of type-b sources outside the box. By [9], the transport capacity for any arbitrary network is no more than $\Theta(\sqrt{n})$, and hence $\tilde{a} \le C\sqrt{n}$. By Equation 9, it follows that for any source outside the $\mathcal{B}_{\frac{\epsilon}{M(n)}}$, the number of lines through tile j can be no more than $C \log n$. Thus, the contribution of type-b sources outside the box is $\Theta(\sqrt{n}\log n)$. By utilizing a scheduling scheme similar to proof of (iii) in Theorem 5.2, each line can be provided a throughput of $\tilde{\Theta}(\frac{1}{\sqrt{n}})$, and hence all the traffic demands of both type-a and type-b nodes are satisfied. ∎

Remark 6.1: The above proof can be generalized to show that any *achievable* rate matrix Λ, with unique source destination pairs ($\Lambda_{i,j} > 0 \rightarrow \Lambda_{i,k} = \Lambda_{r,j} = 0, \forall k \neq j, r \neq i$) can be supported (up to a poly-logarithmic factor) by our algorithm. The RANDOMSPREAD algorithm is modified so that all sources are type-b, and any source i sends $\lceil \sqrt{n} \times \Lambda_{i,d(i)} \rceil$ packets out at any instant. By reformulating the optimization problem of (11) with a_i equal to *the number of packets/lines* generated by type-b sources within $\mathcal{B}_i - \mathcal{B}_{i-1}$, a similar bound on the number of packets through any tile can be shown.

A. Extending Non uniform traffic to Networks with Holes

In Section IV, we demonstrated a randomized routing algorithm to support uniform traffic in the presence of routing holes, and in Section VI a method to support non uniform traffic in uniform random planar networks. Here, we provide an algorithm to combine the two scenarios to provide routing support of non-uniform traffic demands in networks with holes. Consider the following scheme.

Extended RANDOMWAY(n,K): The modification to the RANDOMWAY(n,K) is only at the source nodes. The behavior of the forwarding nodes or the intermediate random-waypoints is unchanged. If the source node is a type-a node, the algorithm is unchanged, i.e, the source creates $R\log(n)$ copies of the single packet and sets the variables in the packet header as described in Section IV. If a source is a type-b node with $\Theta(1)$ traffic requirement, then the source transmits $\frac{\sqrt{n}}{(R\log n)^P}$ packets simultaneously, by executing the RANDOMWAY algorithm for each packet independently, with $4K+2$ way points. That is, for each unique packet p_i, $1 \leq i \leq \frac{\sqrt{n}}{(R\log n)^P}$, the source creates $R\log(n)$ copies and forwards them to random-waypoints.

The above algorithm can be shown to support any *achievable* non-uniform rate matrix, even with network holes. A formal proof is omitted due to space constraints. The algorithm's operation is similar to the RANDOMSPREAD, where nodes with higher traffic requirement send out more packets in proportion to their demands. To facilitate the point-to-point routing (required by RANDOMSPREAD), the intermediate nodes use the RANDOMWAY algorithm to get around possible holes.

VII. Conclusion and Future Work

In this paper, we presented algorithms for throughput optimal routing in networks with holes and non-uniform traffic. Our algorithms preserve the inherent advantages of geographic routing such as scalability and fast convergence while providing better throughput. In future, we will extend the analysis to networks with a larger class of holes and also will characterize the performance under erroneous geographic information.

References

[1] K. N. Amouris, S. Papavassiliou, and M. Li. A position based multi-zone routing protocol for wide area mobile ad-hoc networks. In *Proc. of the 49th IEEE VTC*, pages 1365–1369, 1999.
[2] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, Englewood Cliffs, NJ, 1987.
[3] V. Borkar and P. R. Kumar. Dynamic cesaro-wardrop equilibration in networks. In *IEEE Trans. on Auto. Control, 48(3):382–396*, 2003.
[4] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.
[5] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah. Throughput delay trade-off in wireless networks. In *Proc. of IEEE Infocom*, March 2004.
[6] Q. Fang, J. Gao, and L. Guibas. Locating and bypassing routing holes in sensor networks. In *Proc. of IEEE Infocom*, March, 2004.
[7] Q. Fang, J. Gao, L. J. Guibas, V. deSilva, and L. Zhang. Glider: Gradient landmark-based distributed routing for sensor networks. In *Proc. of IEEE Infocom*, March, 2005.
[8] P. Gupta and P. R. Kumar. A system and traffic dependent adaptive routing algorithm for ad hoc networks. In *Proc. IEEE 36th Conference on Decision and Control*, 1997.
[9] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. on Info. Theory*, IT-46(2):388–404, March 2000.
[10] D.B. Johnson, D.A. Maltz, and J. Broch. DSR: The dynamic source routing protocol for multi-hop wireless ad hoc networks. In *Ad Hoc Networking*, pages 139–172. Addison-Wesley, 2001.
[11] B. Karp and H. T. Kung. GSPR: Greedy perimeter stateless routing for wireless networks. In *Proc. of the ACM/IEEE International Conf. on Mobile Comp. and Networking*, page 243, August 2000.
[12] Y.J Kim, R. Govindan, B. Karp, and S. Shenker. On the pitfalls of geographic face routing. In *Proc. of DIALM-POMC*, pages 34–43, New York, NY, USA, 2005. ACM Press.
[13] E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks. In *Proc. of the 11th Canadian Conf. on Comput. Geom.*, August 1999.
[14] F. Kuhn, R. Wattenhofer, and A. Zollinger. Worst-case optimal and average-case efficient geometric ad-hoc routing. In *Proc. of ACM MobiHoc*, 2003.
[15] T. Leighton and S. Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *FOCS*, 1988.
[16] X. Lin, G. Sharma, R. Mazumdar, and N. Shroff. Degenerate delay-capacity tradeoffs in ad-hoc networks with brownian mobility. In *IEEE Trans. on Info. Theory*, 2006.
[17] R. Madan and D. Shah. Capacity-delay scaling in arbitrary wireless networks. *Allerton Conference on Comm., Cont., and Comp.*, 2005.
[18] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
[19] M.J. Neely and E. Modiano. Capacity and delay tradeoffs for ad hoc mobile networks. In *IEEE Trans Inf. Theory*, June 2005.
[20] C. Peraki and S. D. Servetto. Capacity, stability and flows in large scale random networks. In *Proc. IEEE Information Theory Workshop*, 2004.
[21] C. Perkins and P. Bhagwat. Highly dynamic destination-sequenced distance-vector routing for mobile computers (DSDV). In *Proc. of ACM Sigcomm*, London, UK, August 1994.
[22] C. Perkins, E. Royer, and S. Das. Ad hoc on demand distance vector (AODV) routing. http://www.ietf.org/internet-drafts/draft-ieft-manet-aodv-03.txt, 1999.
[23] E. Peserico and L. Rudolph. Robust network connectivity: when it's the big picture that matters. *SIGMETRICS Perform. Eval. Rev.*, 34(1):299–310, 2006.
[24] M. Shaked and J.G. Shanthikumar. *Stochastic Orders and their Applications*. Academic Press, 1994.
[25] G. Sharma, R. Mazumdar, and N. Shroff. Delay-capacity tradeoffs in mobile ad-hoc networks: A global perspective. In *Proc. of IEEE Infocom*, 2006.
[26] I. Stojmenovic and X. Lin. Gedir: Loop-free location based routing in wireless networks. *Proc. IASTED Int'l Conf. Parallel and Distributed Computing and Systems*, pages 1025–1028, 1999.
[27] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling for maximum throughput in multihop radio networks. *IEEE Trans. on Auto. Control*, 37(12):1936–1949, December 1992.
[28] L. G. Valiant and G. J. Brebner. Universal schemes for parallel communication. In *Proc. of the 13th annual ACM symp. on Theory of computing*, pages 263–277, New York, NY, USA, 1981. ACM Press.