Routing over Multi-hop Wireless Networks with Non-ergodic Mobility

Chris Milling, Sundar Subramanian and Sanjay Shakkottai
Department of ECE
The University of Texas at Austin
Email: {milling,ssubrama,shakkott}@ece.utexas.edu

Randall Berry
Department of EECS,
Northwestern University
Email: rberry@ece.northwestern.edu

***Abstract*—Routing to mobile nodes in a wireless network is conventionally performed by associating a static IP address (or a geographic location) to each node, and routing to that address using routing tables at intermediate nodes that are updated periodically to reflect mobility-induced network topology changes. This mode of routing works when the mobiles' speeds as well as the number of mobiles are small. However, in the presence of large number of fast-moving mobiles, such approaches are infeasible and can lead to excessive overheads, routing failures and hence, throughput loss.**

In this paper, we consider a wireless network over a domain with a collection of static nodes (that form a connected cover of the domain) and mobile nodes, where the mobile nodes can move in an arbitrary (non-ergodic) manner over sub-domains of the network. For such a system, we develop new routing algorithms (based on a spatial multi-resolution search) that we show are efficient both in terms of routing overheads and throughput. In particular, we show that the achievable rate region of the proposed algorithm is within a poly-logarithmic constant of the optimal rate region with non-ergodic mobility.

I. INTRODUCTION

Routing algorithms for MANETs are typically designed by constructing routing tables for the current positions of the mobile destination nodes and periodically updating the tables (proactive protocols) [11] or constructing routing tables in an on-demand manner whenever a packet is required to be routed to the mobile [12]. In the presence of a large number of and/or fast moving mobiles, the routing overheads incurred to learn and maintain correct routing tables at intermediate nodes can be prohibitively large.

Geographic algorithms are advantageous from the perspective of maintaining routing tables as the routing is based only on the locations of the neighboring nodes and the final destination [7]. In recent research [17], [18], geographic routing schemes have been shown to operate close to the optimal rate region (i.e., achieve simultaneous data-rates that are close to rates achievable by any other global algorithm in a static or quasi-static mobility regime).

When the mobiles are slow moving and/or their geographic locations are well known to the sending node, it is possible to construct the routes and routing tables with low-overhead costs and amortize the associated route setup costs over multiple data packets. However when the locations of the mobiles are unknown to the sending node (and the nodes are fast moving), the overhead required for route setup can become significant, precluding the use of conventional geographic algorithms.

Another important consideration in designing algorithms for mobile routing is the mobility model. In many studies [4], [13], the pattern of mobility is assumed to be ergodic, i.e., the mobiles are assumed to move with some statistical regularity over a given set/region. Such an assumption implies that a mobile will return to fixed locations in the region with probability one, enabling the development of algorithms for near-throughput-optimal routing [14], [19]. However, such an assumption might not be valid in many applications.

In this paper, we consider a scenario where a collection of static (non-mobile) nodes form a (wireless) connected cover of the spatial domain, and a large number of (potentially fast-moving) mobile nodes move over this region in a potentially non-ergodic manner (i.e., there is no guarantee that the mobiles' trajectories have any learnable statistical regularity). Our objective is to develop multi-hop routing algorithms to route data from the static nodes to the mobile nodes assuming that the static nodes are not aware of the mobile nodes' locations. However, the static nodes have local geographic knowledge (i.e., information about their neighbors within their radio range as well as their own locations) to enable geographic forwarding. Such a scenario has several applications – for instance, routing from a collection of sensors nodes scattered over a battlefield to mobile soldiers [14], or industrial applications where sensory data from different locations need to reach (mobile) foremen. In either case, it is not clear that we can precisely characterize the mobility pattern beyond some rough measure (such as approximate regions over which the mobile can move).

In this regime, we develop a mobile-assisted sequential-search routing algorithm that performs near to the optimal rate region of networks with arbitrary mobility over sub-regions of the geographic domain [9].

A. Main Contributions

We consider a network where n static nodes are distributed uniformly randomly over a unit radius sphere, and there are n mobile nodes that serve as destinations. Associated with each mobile $i, 1 \leq i \leq n$ is a region of the surface of the unit sphere, $\mathcal{S}_i$, over which mobile i can move in an non-ergodic manner. The mobile nodes are also capable of sending out advertisement packets to announce their current locations. Further, associated with each mobile i is a static node that generates traffic to that mobile at rate λ_i. We study the performance of such a network as n scales assuming that the static node associated with mobile i is $\Theta(1)$ distance away from the mobility set $\mathcal{S}_i$[1]. For such a system, our contributions are as follows:

1) We construct a routing algorithm based on a mobile-assisted sequential search that iteratively gets closer to the mobile destination irrespective of its mobility. The routing length (hop-count) due this scheme is shown to be order-wise the same as a direct path from the source to the mobile.
2) The algorithm is analytically shown to be correct, and to achieve a rate region that is within a poly-logarithmic fraction of an upper-bound on the optimal rate region with non-ergodic mobility i.e., the algorithm is shown to be 'near-optimal' in terms of the rate region.

B. Related Work

Many routing algorithms for MANETs have been extended from existing algorithms for static wireless networks like DSDV [11], AODV [12], and DSR [6]. The means of using a static algorithm is either by updating the routing table whenever the network topology is altered due to mobility (proactive protocols) or by flooding the network to construct a path whenever a packet needs to be routed (reactive protocols). With both approaches, the routing overheads (table updates and/or reactive flooding) increase considerably with the rate of change of network topology. In networks with a small number of slow moving mobiles, the overhead costs of such approaches could be amortized over multiple packet transmissions. However, in networks with a large number of fast moving nodes, such topology changes can be very frequent and cause the routing schemes to fail or perform poorly.

To ameliorate the overhead cost due to topology changes, algorithms based on geographic forwarding have been studied as a popular alternative [7], [1]. With geographic forwarding, as the routes can be constructed via greedy forwarding, the nodes only need to update the locations of their neighbors. However, these schemes require the source nodes to know the destination's location. Under such assumptions, recent work [17], [18] has shown that variations of geographic routing with appropriate randomization can achieve 'near optimal' throughput with arbitrary traffic distributions and network holes (regions where geographic forwarding may fail). However, when the destination is a fast moving mobile node, any location information available at the source will quickly become outdated, causing geographic schemes to also perform poorly or fail.

The idea of reducing (making infrequent) the topology updates to nodes far away has been considered in schemes such as FISHEYE state routing [10], which provides increasingly precise route updates as a packet approaches the mobile destination. Routing with mobility has also been considered in recent research where the movement of the mobile nodes has mainly been seen as a means to increase throughput. Authors in [4] consider a network in which mobiles move ergodically over the whole region and show that a throughput capacity (i.e., the maximum data-rate that is simultaneously achievable by all source-destination pairs) of $\Theta(1)$ is achievable. This capacity is mainly achieved by using the mobiles to relay packets to other mobiles whenever they come close enough to each other (this happens roughly periodically due to ergodicity). A critical aspect of using the mobiles to relay the packets is the associated increase in the delay as the time taken by a relay mobile node to revisit a given mobile is much larger than the time needed to route a packet directly to that node. The throughput-delay trade-off in networks with mobility (ergodic mobility over the whole region or random walks over multiple intersecting regions that allow for sufficient mixing) has been well studied in [3], [15], [8], [13], [19].

Mobility over restricted sets (but ergodic) was considered in [14] where the authors demonstrated that simple geographic schemes can perform near-optimally (in a throughput capacity sense) by routing packets to a random node in the mobile's restricted set. We note however that in these approaches, the mobility is ergodic, i.e. nodes are visited with regularity, and hence packets can wait for mobiles to collect them.

In the absence of any ergodic movement (i.e., any learnable mobility statistics/pattern) the above schemes perform poorly with respect to capacity and could even fail to route packets successfully. The regime of non-ergodic mobility and a characterization of the throughput region (rate region) was presented in [9] in a cellular context.

In this paper, we consider the non-ergodic movement of mobiles over a connected set $\mathcal{S}_i$ and demonstrate a routing algorithm based on using geographic forwarding and advertisement by mobiles to sequentially partition the network topology so that packets reach their destinations over multiple iterations. The use of advertisements by mobiles to aid routing has also been previously considered [2], but in the context of querying and searching sensor networks.

II. System Description

A. Networks with Arbitrary Mobility

We consider a network where n static nodes are distributed uniformly randomly over a sphere with unit radius. Further, there are n mobile nodes located arbitrarily. These mobile nodes move non-ergodically (arbitrarily) through some region $\mathcal{S}_i$ for each mobile i in the network. Each mobile node can receive packets from nearby static nodes (within radio range) at rates much higher than the static node-to-node rates. Thus we assume that a mobile picks up all packets destined to it in a single time unit whenever it visits a static node, see Figure 1, [14]. We do not allow mobiles to carry packets for a different mobile. Thus node-to-mobile communication is limited to node-to-node hops until the last transmission to the correct mobile. The data rates are such that at least some fraction of the rate to any particular mobile comes from a

[1]This assumption ensures that the rate a static node can send to a mobile node is bounded.

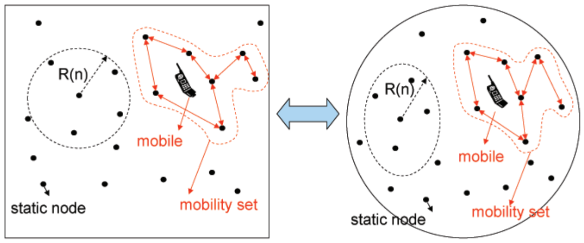

Fig. 1. System model illustrating static and mobile nodes with restricted mobility sets.

Field Name	Functionality
ADVERTISEMENT or DATA	Mobile information or Data Packet.
ROUTING STAGE	Specifies the current stage of routing
ACK	Denote packet receipt.
CURRENT CENTRAL NODE	Central node, random grid orientation
ITERATION	The current iteration level
SEC-DEST	Location of next+1 way-point
MOBILE-DEST	ID of the mobile destination
DATA	Message to the destination node
TIMESTAMP	Timestamp for advert./control packet

TABLE I
FIELDS IN THE HEADER OF THE PACKET.

source at a distance $\Theta(1)$. That is, we only allow rates such that at least f_{min} fraction of the rate comes from a distance which is[2] at least r_0.

The authors in [5] have shown that for any $\epsilon > 0$, one can construct a Voronoi tessellation of the surface of the unit sphere so that each region can be bounded inside and outside by circles of radii ϵ and 2ϵ respectively. We tile the spherical region by such a tessellation with $\epsilon = \Theta(\sqrt{\log n/n})$. The radio range $R(n)$ is also chosen to scale as $R(n) = \Theta(\sqrt{\log(n)/n})$, and each static node can communicate to nodes in adjacent tiles. By the setting appropriate constants for ϵ and the radio range (i.e., the constants in the $\Theta(\cdot)$ term), the connectivity of the network can be guaranteed almost surely (as $n \to \infty$) [5]. Henceforth in this paper, we assume that n is large enough such that geographic routing will not fail with high probability. Also, there are no more than $\Theta(n/\log n)$ tiles since each tile covers at least a $\Theta(\log n/n)$ area.

B. Interference Model and Standard Definitions

We assume the following to model the interference effects of simultaneously transmitting nodes which are within each other's radio range $R(n)$.

Definition 2.1 (Protocol Model, [5]): A transmission between a node A and its receiving node B is assumed to be successful if $d(A,B) \leq R(n)$ and $d(C,B) > (1+d)R(n)$, for some $d > 0$, for all other transmitting nodes $C \neq A$.

This successful transmission occurs at rate '1' WLOG. Further, we denote by the *rate region*, the set of all n dimensional achievable rate vectors. Thus, each n dimensional rate vector in the rate region corresponds to the data rate that can be simultaneously sustained between the n source-destination pairs under some (possibly mobility-pattern aware) routing strategy, and with any arbitrary mobility pattern as described in Section II-A. More formally, let $\bar{\mathbf{\Lambda}}_{mp,s}$ be the set of n dimensional rate vectors that can be achieved under a routing strategy s and a mobility pattern mp (where this mobility pattern is feasible under the constraints imposed in Section II-A). The the rate region $\bar{\mathbf{\Lambda}}$ is given by

$$\bar{\mathbf{\Lambda}} = \bigcap_{mp} \bigcup_{s} \bar{\mathbf{\Lambda}}_{mp,s} \tag{1}$$

Note that in the above definition, the routing strategy could adapt according to the mobility pattern (i.e., mobility aware routing). 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 $-\infty < k, k_1 < \infty$. When applied to collections of scalars (e.g. vectors, matrices, sets), the $\tilde{\Theta}(\cdot)$ notation applies for each scalar component of the collection. Thus, if $\bar{\mathbf{\Lambda}}$ is the rate region, we say that an algorithm is near-optimal if it achieves a rate region that is $\tilde{\Theta}(\bar{\mathbf{\Lambda}})$.

III. ROUTING TO MOBILE NODES

Routing to static nodes in a wireless network can be performed very efficiently in the presence of geographic location information [7], especially when the location of the destination node is known to the sender. With geographic location, the packets can be greedily forwarded towards the respective destinations. When the destinations are mobile, (and/or their locations are unknown) such an approach is infeasible.

In scenarios where the mobile speeds are much slower than packet speeds (distances traveled by a typical packet in a given time), a possible routing method maybe to chase down a traveling mobile. Also, when the movement of the mobiles is ergodic (i.e. mobiles revisit given nodes with high probability), a possible routing scheme is to route a packet to one of the nodes that the mobile shall eventually visit, and hold the packet until the mobile arrives. However, both of these methods could fail when the mobile's trajectory is arbitrary and fast moving.

In this section, we describe a routing algorithm that successfully delivers the packets to the mobile irrespective of its speed. Also, the algorithm is shown to provide 'good' throughput even in the presence of arbitrary mobility, i.e, the scheme achieves a rate region that is within a poly-log factor of the best possible rate region (over the worst-case mobile scenarios as described in Section II-B).

A. MobileSearch Algorithm

We first define a packet structure in Table I to provide a common communication scheme between nodes. The routing process is described as a logarithmic search for the mobile

[2]Note that this is in fact a weaker requirement than having a single static source node associated with every mobile. This condition allows multiple source nodes to transmit to a mobile destination – however, in the rest of the paper, we state and prove results only in the context of one static source per mobile.

node over the network. For each packet, the transmitting node constructs a random grid (i.i.d. for each packet) that sequentially partitions the whole network space into finer regions until the packet reaches the mobile. We initially describe the algorithm over a unit square region for ease of description, although in the analysis we use the correspondingly created grids on a unit sphere, see Figure 2.

The source node A sets the starting iteration number $i = 0$ and picks a randomly chosen point P along with a random orientation; it then forwards the packet to the node at position P and indicates the orientation and iteration in the header. Note that on a sphere, the point P is uniformly spatially symmetric. The randomly chosen central node P then constructs the grid lines G, G' and H, H' to divide the whole region into four sub-pieces (i.e., it quarters the region). To do this it sends the packet over each grid line[3]. Each node along the grid saves a copy of the packet for a predetermined time. This can be performed in such a way that the packet passes along each line segment at most twice and returns to the central node P. Since the static nodes are aware of their own geo-location, this traversal can be performed by greedy packet forwarding where each node needs only to know the orientation and level of the grid, and the end points of the current line segment to determine how the packet can be forwarded. At this stage, we designate the whole region as the 'active-grid', and the grid-lines as the 'current' grid lines. The algorithm's operation until this point is designated as the BUILD stage. Once the traversal for constructing the 'current' grid is complete, the central node initiates the WAIT stage. First, a message packet specifying the stage is sent along the grid in the same manner as described above. Once nodes receive and forward this message, they simply wait in an inactive state.

An important component of the algorithm is the advertisement by the mobile destinations. Each of the mobiles periodically send out advertisement messages along a line towards a random point[4]. This message will advertise the mobile's position with the time it was present in that location. Eventually, a node along the 'current' grid will receive this message when it is in the WAIT stage. On receiving the advertisement, the position contained in the message is forwarded along the grid to the current central node. Further, the advertisement packets are timestamped. The advertisement will be used to activate a new 'active-grid' only if the BUILD completion time is earlier than the advertisement timestamp (also note that multiple advertisements in an iteration will be suppressed by the grid nodes).

The central node uses the position of the mobile to determine the section of the grid where the mobile is present. A new central node P is chosen in the section containing the mobile, and the packet is sent along the next level of the grid to that central node. Note that the header will still contain the information about the grid's orientation. At this stage, the iteration number is incremented by one. Once the new central node receives the packet, the process is repeated at the smaller sub-region. Now, the 'active-grid' is the sub-region where the mobile was present, and the 'current grid' is comprised of all the newly constructed finer grid-lines along with the boundaries of the 'active' region (see Figure 2).

[3]In the following when we say the source sends a packet over a grid line, we mean that it geographically forwards the packet along the static nodes in the direction of the grid line.

[4]Note the random point establishes the direction of the line, but is not the end point. The packet will continue in this direction until it intersects the grid or reaches the boundary or the region (in case of the sphere, completes the great circle).

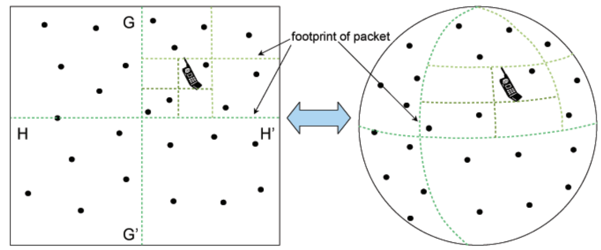

Fig. 2. Construction of the sequentially refined grids to trap the mobile node. This figure illustrates the algorithm on a unit planar area as well as on a unit sphere.

This packet is successfully received in one of two ways. Whenever the mobile moves over a tile, it picks up all the packets within that cell that are destined for the mobile. Thus if the mobile moves outside of the grid enclosing it, the packet is successfully received. Otherwise, the subdivision continues until the grid contains only some small constant number of cells[5]. The packet is then just spread over this small area and the mobile must receive it. Once a packet is received by its destination an acknowledgment is sent to the appropriate central node and the routing of the packet will terminate within one iteration of the algorithm.

B. Analysis of MobileSearch Algorithm

In this section, we demonstrate the correctness of the algorithm in routing the packets to their corresponding mobiles and its near-optimal throughput property, i.e., the algorithm successfully achieves a poly-log fraction of the rate region. We first provide an upper bound on the achievable rate region in the non-ergodic mobility regime. We then show that the mobile is always contained in the correct 'active-grid' at any given iteration of routing a packet, irrespective of the mobile's movement. Finally, we prove that the network will be stable for rates that are in a poly-log fraction of our bound on the rate region. That is, we consider a typical tile of the dimension of the radio range, and show that the rate imposed by our routing algorithm does not overload any given tile - the rate imposed is within a poly-log multiple of the maximum rate (a $\Theta(\sqrt{n})$ quantity over time-slots $\sqrt{n}$ time-units long) the tile can support.

As mentioned in Section II, we consider a Voronoi tessellation of the sphere such that the space is tiled by approximately uniform tiles of the dimensions of the radio

[5]Note this implies that the maximum number of iterations will be $\Theta(\log n)$.

range $R(n)$. Recall that the radio range is chosen to scale as $R(n) = \Theta(\sqrt{\log(n)/n})$ to ensure that the static network is well connected with high probability for large values of n [5]. We choose a typical tile from such a construction. We assume that nodes send constant-sized packets at a maximum transmission rate of 1 packet per time-unit from each cell (independent of n). However, actual packets are sent over a time-slot that is $\sqrt{n}$ time-units long, which enables up to $\sqrt{n}$ packets to be multiplexed in each time-slot from a given tile. For example, a node that wishes to communicate at a rate $1/\sqrt{n}$ would send out one packet each time-slot. Nodes that wish to communicate at a higher rate would send more packets. We comment that an alternate approach is to construct fixed-size time-slots and scale the packet size as $1/\sqrt{n}$ [3], which will lead to the same packet data-rate as the scaling we use in this paper. However, as we have various messages relayed as part of the protocol (e.g., advertisement packets, ACK packets) and it is not clear how each of them should scale, we assume fixed size packets that do not scale with network size, and instead scale the time-slot length as $\sqrt{n}$.

We now provide an upper bound on the rate regions that can be supported with non-ergodic mobility. Such an upper-bound has also been provided in [9]. We provide the theorem in the context of our system and a short proof for completeness.

Theorem 3.1: Let $\bar{\mathbf{\Lambda}}$ be the n-dimensional rate region for a system of n mobiles with non-ergodic mobility as described in Section II-A. Then,

$$\bar{\mathbf{\Lambda}} \subset \cap_{S \in \mathfrak{S}} \tilde{\bar{\Lambda}}_S, \tag{2}$$

where $\mathfrak{S}$ is the set of all allowable static configurations of the mobiles allowable under the constraints imposed in Section II-A, and $\tilde{\bar{\Lambda}}_S$ is the collection of rate vectors that are supportable for the static configuration S.

Proof: The proof follows since non-ergodic mobility allows for nodes to also be static in any allowable position. Thus, the capacity region can be no greater than the intersection of the rate regions achievable by static configurations. ■

The following lemma provides a proof that our routing algorithm (i.e, the per-packet iterative search for a mobile) correctly traps the mobile in its current ‘active-grid’ (see Section III-A) at every stage of routing. From the discussion in Section II, we have that the static network is well connected and geographic routing (i.e., along straight-lines) will not fail with high probability as $n \to \infty$. Thus, to simplify notation and exposition, in the subsequent theorems, we assume that the static network is densely connected so that geographic routing will not fail.

Theorem 3.2: The ‘active-grid’ $H_{i,p}$ constructed at iteration i for routing the unreceived packet p to a mobile node M contains the mobile M.

Proof: Let p be some packet that is sent to a mobile node M. We will prove this theorem using induction over the iterations of the algorithm.

Let E be the number of iterations of the algorithm. We will show that at the start of any iteration i, $0 \leq i \leq E$, either the packet p was already successfully received, or the mobile M is inside $H_{i,p}$, the ’active-grid’ for iteration i. The first grid is over the entire network, so clearly the mobile must be contained in it. This covers the base case, as $H_{0,p}$ always contain M.

For the inductive step, assume that the induction hypothesis holds for iteration k, where $0 \leq k < E$. Namely, assume that either p was received before the start of iteration k, or that M is inside the region $H_{k,p}$ at that time. We will show that from this, we can conclude that at the start of the next iteration $k+1$, p was received or that M is inside $H_{k+1,p}$. If the first condition of the induction hypothesis is true, i.e. that p was received before iteration k, then clearly p was received before iteration $k+1$ and we are done. Otherwise, the mobile is inside the region $H_{k,p}$ at the start of iteration k. During that iteration, the first stage of the algorithm is the BUILD step. In this step, this grid is subdivided into four sections, one of which will become the next iteration’s ’active-grid’. Then the algorithm proceeds into the WAIT stage.

There are two possible cases at this point. Either the mobile M left the grid $H_{k,p}$ or it has remained there. In the former case, where the mobile M is no longer in the grid $H_{k,p}$ when the algorithm proceeds to the WAIT stage during the k^{th} iteration, clearly it must be elsewhere in the network. By the structure of the grid, we know that all nodes along the border of the ’active-grid’ $H_{k,p}$ contain the packet p. Note that there is no border around $H_{0,p}$ (since it is the whole network), but the mobile cannot travel outside that region. Since the path of the mobile is continuous and travels from inside the region $H_{k,p}$ to outside of it, it must pass through the border. Thus M must have passed within range of a node containing p. By assumption, when the mobile passes within radio range of a node, the mobile is capable of receiving all packets meant for it from that node. Therefore, M would have received packet p at this time, before the start of iteration $k+1$.

In the second case, the mobile is still inside $H_{k,p}$ when the algorithm enters the WAIT stage. Note that at this time, the BUILD stage is complete and the packet p has been sent over each of the lines in the next finer grid. Recall that the mobile M is periodically sending out advertising lines. Then at some point after this WAIT stage starts, M will send out one of these advertising line that will intersect the grid. The node at this intersection point uses the timestamp inside the packet to verify that it was sent after this stage began. The packet gives the location of the mobile M. Like before, if M left $H_{k,p}$ before sending the line, it received p. Otherwise, M sent the line while inside one of the sections of $H_{k,p}$. This section will become the ’active-grid’ for the next iteration, $H_{k+1,p}$. Finally, when the $k+1$st iteration begins, either the mobile is still in $H_{k+1,p}$ or it left. As mentioned earlier, p was already sent to all nodes bordering $H_{k+1,p}$ (and each other section of $H_{k,p}$) during the BUILD stage. Hence, using continuity of motion as before, M must have passed near a node with the packet p on the border of $H_{k+1,p}$, receiving the packet. By combining all the cases, we have that either the mobile M passed through the grid and thereby received the packet p, or it is inside the $H_{k+1,p}$ at the start of iteration $k+1$.

Continuing the induction to the last iteration of the algorithm, we find that at each iteration, the packet was already received or the mobile M is inside the 'active-grid' for that iteration. ■

Next, we provide bounds on the probability that the routing grid constructed by our algorithm for a mobile M can enter into a typical tile T, which is at a distance $r = d(M,T)$ from the mobile. For any given packet p, we define the set A_p as the *footprint of routing - i.e., set of tiles touched by the routing scheme.* Note that our routing adds new tiles to the footprint at every iteration i, denoted $A_{i,p}$. Thus $A_p = \bigcup_i A_{i,p}$ (the union is over only those iterations until the algorithm terminates).

Lemma 3.1: Let M be a mobile receiver, T be some typical tile and let $A_{i,p}$ be the set of new tiles added to the footprint at iteration i. Then, if $d(M,T) > \Theta(2^{-i})$, the probability that $T \in A_{i,p}$ is 0. Otherwise, $\mathbb{P}\{T \in A_{i,p}\}$ is at most $\tilde{\Theta}(2^i/\sqrt{n})$.

Proof: We need only consider when the algorithm reaches iteration i (because the packet p was not yet received), or otherwise $A_{i,p} = \emptyset$ and the upper bounds trivially hold. Let d_{max} be the maximum diameter of the 'active-grid' for p at iteration i, $H_{i,p}$, over all possible regions. Note d_{max} is a $\Theta(2^{-i})$ quantity since the sides of each region are cut approximately in half each iteration. Note that all tiles added to the footprint at iteration i must at least partially intersect $H_{i,p}$.

Consider the case where $d(M,T) > d_{max}$. By Theorem 3.2, we know that $M \in H_{i,p}$ at this iteration. Suppose it was possible that $T \in A_{i,p}$. Then some part of T must intersect $H_{i,p}$. The distance between any point in $H_{i,p}$ and the tile T is at most d_{max} since the tile T intersects with $H_{i,p}$ whose diameter is at most d_{max}. Yet $d(M,T) > d_{max}$ and $M \in H_{i,p}$. This is a contradiction. Hence if $d(M,T) > d_{max}$, it is impossible for $T \in A_{i,p}$, and the probability that this event occurs is 0. This completes the proof of the first part of the lemma.

Next, we would like to bound the desired probability in the case where $d(M,T) \le d_{max}$. Here, we ignore the 'active-grid' and consider the extension of the 'active-grid' over the whole network. That is, we consider the grid that is the union of the lines in all possible active-grids from any mobile position (the actual 'active-grid' that is used depends on the mobile position). This is a grid at the resolution of iteration i that covers the whole space. We call the lines making up this grid G_i. Note that each iteration at most doubles the number of lines in the grid, each of which has a maximum length. This implies that the total length of the lines in G_i is at most $\Theta(2^i)$. The packet actually travels along a subset of this grid that is within the active area. Let $B_{i,p}$ denote the footprint from a packet traveling over the extended grid G_i, i.e. the set of tiles that intersect with G_i. Since the packet travels a subset of this grid, $A_{i,p} \subset B_{i,p}$. This gives the following upper bound: $\mathbb{P}\{T \in A_{i,p}\} \le \mathbb{P}\{T \in B_{i,p}\}$.

Note that $B_{i,p}$ is independent of the location of the mobile. Recall 2ϵ is the maximum radius of the tile T, where $\epsilon = \Theta(\sqrt{\log n/n})$. Also note that any grid can be traversed by a path that crosses any point at most twice. Let P_i be such a path that traverses G_i, so $\ell(P_i) \le 2\ell(G_i) = \Theta(2^i)$ ($\ell(\cdot)$ is the length of the corresponding path). Mark out D points labeled $\{S_0, S_1, \ldots, S_D\}$ over the whole length of P_i at intervals of ϵ, including the first and last points of P_i. Now, noting that $D \times \epsilon/2 < \Theta(2^i)$, we have $D \le \tilde{\Theta}(2^i\sqrt{n})$.

We call the center of the 2ϵ radius circle that covers T the center of T, and label it T_0. Consider a circle C of radius 4ϵ around T_0. We will upper bound the probability that $T \in B_{i,p}$ by considering the probability of a simpler event $T \cap G_i \ne \emptyset$. Suppose there is some point $p \in G_i \cap T$. It must lie in some interval $\overline{S_j S_{j+1}}$ on the path P_i. Note that due to the construction of that interval, $d(p, S_j) < \epsilon$. The triangle inequality gives $d(S_j, T_0) \le d(S_j, p) + d(p, T_0) < 3\epsilon$. Thus, point S_j is in the circle C. Then we have that $G_i \cap T \ne \emptyset$ only if there is at least one of the points S_j inside C. Hence,

$$\begin{aligned} \mathbb{P}\{T \in A_{i,p}\} &\le \mathbb{P}\{T \in B_{i,p}\} \le \mathbb{P}\{\exists j : S_j \in C\} \\ &\le_{(a)} \sum_{j=0}^{D} \mathbb{P}\{S_j \in C\} =_{(b)} \sum_{j=0}^{D} Area(C)/(4\pi) \\ &=_{(c)} (D+1)\tilde{\Theta}(\frac{1}{(\sqrt{n})^2}) \le \tilde{\Theta}(2^i/\sqrt{n}) \end{aligned}$$

Inequality (a) above follows from a union bound. The next equality (b) follows due to the randomization of the grid. Any point S_j on the grid is uniformly probable to be any point on the sphere. Then the probability that it is in a given circle is the area of that circle divided by the surface area of the whole sphere, 4π in this case. Using the radius of C gives the equality (c). Thus we have $\mathbb{P}\{T \in A_{i,p}\} \le \tilde{\Theta}(2^i/\sqrt{n})$, which completes the lemma. ■

While Lemma 3.1 provided a probabilistic bound on the rate at which packets may enter a tile given the mobile's distance from it, the following lemma provides a cut-set based upper bound on the sum-rate at which packets (and the number of mobiles) can enter a radial region.

Lemma 3.2: For any rate vector in the rate region, the following property holds: In any circle with radius r, with $2r < r_0$, the total rate destined for any mobiles inside that circle is at most $\tilde{\Theta}(rn)$.

Proof: Let A_r be the set of all mobiles that may be in the circle. This equals all mobiles whose mobility set intersects with A_r. Note that there is a static configuration in which each of these mobiles are in the circle, and hence by Theorem 3.1, the optimal capacity region is a subset of the rate region for this configuration. This result is independent of mobility and gives a bound on the rates as if all the mobiles were static in the worst case positions for a given traffic vector. For $M \in A_r$, let λ_M be the total rate (in packets per time-slot) that packets are being sent to mobile M. Then by assumption, at least a f_{min} fraction of the rate is from at least r_0 distance away from the mobile. Since the diameter of the circle is smaller than this distance, M must receive a rate at least $f_{min} \times \lambda_M$ from nodes outside the circle. The perimeter of the circle is slightly less than $2\pi r$ (since this is on a sphere). Then a cut-set bound gives the maximum amount of rate that can enter the circle. Only one node can transmit over each $R(n) = \tilde{\Theta}(\frac{1}{\sqrt{n}})$ length of the perimeter and each node can only transmit at a rate $\Theta(\sqrt{n})$. Hence, packets can cross the boundary of the circle at a rate no more than $\tilde{\Theta}(2\pi rn)$. This gives the following bound.

$$\sum_{M\in A_r} f_{min} \times \lambda_M < \tilde{\Theta}(rn)$$

As f_{min} is a $\Theta(1)$ quantity, we have proved the desired result. ■

Theorem 3.3: Let Λ be an achievable traffic vector. Then, Algorithm MobileSearch achieves a traffic vector that is $\tilde{\Theta}(\Lambda)$.

Proof of Theorem 3.3:

To prove this theorem, we will show that each tile of the network receives a load at no more than a $\tilde{\Theta}(\sqrt{n})$ rate in any time slot (which is $\sqrt{n}$ time-steps long). This rate is achievable by a schedule that operates FIFO and tiles are scheduled by a finite coloring [5]. Note that the total traffic can be decomposed into the following three parts.

- Star shaped traffic ($*$-traffic) from the radially outgoing packets at the first leg of routing.
- Traffic generated by the mobile node due to its radially outward advertising line.
- The traffic generated by the sequential grids that have been constructed.

We prove that each of these traffic types contribute no more than $\tilde{\Theta}(\sqrt{n})$ into any tile with high probability.

1) $*$-traffic from Routing: The $*$-traffic comes from packets being sent out to the first central node, which is uniformly chosen over the whole sphere. The traffic here is generated by nodes sending out packets on random lines at a rate equal to the total rate they want to send to mobiles. This type of traffic was analyzed in [18], where it was shown that this traffic produces a load of no more than $\tilde{\Theta}(\sqrt{n})$ into any tile (however, in the context of a torus). A similar proof holds for the unit sphere in this paper; we skip the details for brevity.

2) $*$-traffic from Advertisement: The second type of load is from the advertising lines that each mobile sends. It can be shown that by sending advertising lines proportional to the rate a mobile receives, the load on any tile is also no more than $\tilde{\Theta}(\sqrt{n})$. The proof of this is similar to the $*$-traffic from routing, as the cut-set bound in Lemma 3.2 applies equally to receiving and sending packets.

3) Traffic from the sequential grids: To prove that the traffic from the grids does not overload the tiles, we show the following intermediate steps. Firstly, we show that any tile on the footprint of a packet receives no more than $\Theta(\log n)$ packets of information over a time-slot. Next, we provide a probabilistic bound on the event that a given tile lies in the footprint of a packet. Finally, we construct a series of annular regions (whose radii decrease geometrically) around a typical tile and consider the rate from each of the annular regions that impact the tile. We show that the contribution from each ring (annular region) is $\tilde{\Theta}(\sqrt{n})$ over each time-slot and the number of such annular regions is also logarithmic (see Figure 3). The details are as follows.

Claim 1: Any tile on the footprint of a packet receives no more than $\Theta(\log n)$ packets of information due to that packet.

Proof: For some packet p, consider the tiles in the footprint A_p. Consider a tile T in this footprint and an iteration i in which T receives packets. We will now step through this iteration to determine how many times T may receive information packets. In the BUILD stage, the grid is constructed with the packet traversing any tile in the footprint of the grid up to at most two times. In the WAIT stage, the central node starts by sending a message over each line in the grid to acknowledge that the grid was completely constructed. Again, this is done with that message passing over any tile at most twice. Once the advertising line is received by a node, it forwards an ACK for the advertisement to the central node. Note any node will only forward such a ACK once, and later equivalent ACKs are ignored. Then this adds at most one packet load to the line. Finally, the last ITERATION stage involves send a message along the grid to the next central node. Again, this is a load of at most one. Totaling up the load from each stage, we find that T receives no more than 6 packets during a single iteration. In the worst case, T may receive packets at each iteration, and there are $\Theta(\log n)$ iterations. From this, we find that T receives at most $\Theta(\log n)$ packets of information total. ■

Claim 2: The probability that the footprint of a packet p for mobile M contains a tile T is $\tilde{\Theta}(\frac{1}{r\sqrt{n}})$, where $r = d(M,T)$ at the last position of M before the packet is received.

Proof: Recall from Lemma 3.1, for the footprint from iteration i:

$$\mathbb{P}\{T \in A_{i,p}\} \leq \begin{cases} \tilde{\Theta}(2^i/\sqrt{n}) & r \leq \Theta(2^{-i}) \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

At iteration $F+1$, $\Theta(2^{-(F+1)})$ will decay to less than r (or r will be less than the radio range). At this point, we know T cannot be added to the footprint. For clarity, let c_1 and c_2 be the constants for the $\Theta(2^{-i})$ and $\tilde{\Theta}(2^i/\sqrt{n})$ expressions above. Then F can be determined as follows:

$$\begin{aligned} d(M,T) = r &> \Theta(2^{-(F+1)}) \\ &= c_1 \times 2^{-(F+1)} \\ F+1 &> \log_2(c_1/r) \\ F &\leq \log_2(c_1/r). \end{aligned}$$

Using this, we can apply a union bound to the result of Lemma 3.1.

$$\begin{aligned} \mathbb{P}\{T \in A_p\} &= \mathbb{P}\{\exists i \leq F : T \in A_{i,p}\} \\ &\leq \sum_{i=0}^{F} \mathbb{P}\{T \in A_{i,p}\} < \sum_{i=0}^{F} c_2 \times \frac{2^i}{\sqrt{n}} \\ &< \frac{c_2}{\sqrt{n}} \times \sum_{i=0}^{F} 2^i < \frac{c_2 \times 2^{F+1}}{\sqrt{n}} \\ &< \frac{2c_2 \times 2^{log_2(c_1/r)}}{\sqrt{n}} < \frac{2c_2c_1}{r\sqrt{n}} \\ &= \tilde{\Theta}(\frac{1}{r\sqrt{n}}). \end{aligned}$$

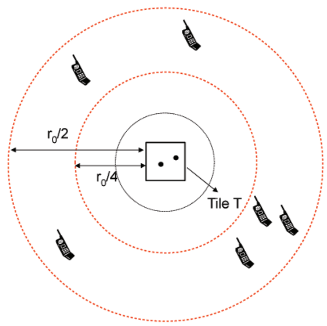

Fig. 3. A concentric collection of annular regions provide mobility independent cut-set bounds on the data-rate that can enter a region, and can be used to compute an upper-bound on the load on tile T.

This is the desired bound. ■

We will now partition the mobiles into several groups. Each successive group will be closer to the tile T and be more likely to add load to T. This will be balanced out by tighter cut-set bounds that will increasingly limit the total rate mobiles in that group can receive packets.

Let U_0 be the set of all mobiles that are at least $r_0/2$ away from T. Since each node can transmit at most $\sqrt{n}$ packets per time-slot, the total number of packets transmitted over the whole network is less than $n\sqrt{n}$ per time-slot. Since we assume that a significant fraction of the packets must travel a $\Theta(1)$ distance and these packets must take $\Theta(\sqrt{n/\log n})$ hops, the load on the network is added at a rate at least $\Theta(\sqrt{n/\log n} \times \mu)$, where μ is the total rate packets are sent. Using the above transport capacity bound, $\Theta(\sqrt{n/\log n} \times \lambda) \leq \Theta(n\sqrt{n})$ so $\mu \leq \tilde{\Theta}(n)$. Note for each mobile in this set, $d(M,T) > r_0/2 = \Theta(1)$ (and $\mu \leq \tilde{\Theta}(\frac{r_0}{2}n)$).

Next we will divide the remaining area into a series of annuli and a circle smaller than the radio range around tile T (see Figure 3). Let U_1 be set of mobiles between $r_0/4$ and $r_0/2$ from tile T, U_2 be the set of mobiles between $r_0/8$ and $r_0/4$ from tile T, and so on. Note here that we mean the set of mobiles that may be in that region at some time. Let r_i denote the inner radius of U_i. Let U_F be the last annulus with $r_F < 2\epsilon = \tilde{\Theta}(1/\sqrt{n})$. That is, $F > \log\sqrt{n} \times 2/r_0 = \tilde{\Theta}(1)$.

Next, each mobile $M \in U_i$ may be between r_i and $2r_i$ away from T. Since M is inside the circle of radius $2r_i < r_0/2$, we know that $\sum_{M\in U_i} \lambda_M < \tilde{\Theta}(r_i n)$ by Lemma 3.2.

Claim 3: Choose any $U \in \{U_0, U_1, \dots, U_F\}$ that is at least r distance from T and satisfies the following rate bound: $\sum_{M\in U} \lambda_M < \tilde{\Theta}(rn)$.

Then during a single time-slot, the load on T due to the mobiles in U is less than $\tilde{\Theta}(\sqrt{n})$ with probability at least $1 - \frac{1}{n^3}$.

Proof: Let X be a random variable equal to load from mobiles in U on T over one time-slot, and $X_M(p)$ be the load from a packet p for a mobile $M \in U$. From Claim 2, for all packets p to a mobile $M \in U$, $\mathbb{P}\{T \in A_p\} = \tilde{\Theta}(\frac{1}{r\sqrt{n}})$. In addition, each packet for mobile M adds a load of at most $\Theta(\log n)$ by Claim 1. Then each $X_M(p)$ is stochastically dominated by Bernoulli random variables $\Theta(\log n) \times \tilde{X}_M(p)$ (independent across packets and mobiles), where

$$\tilde{X}_M(p) = \begin{cases} 1 & \text{w.p. } \tilde{\Theta}(\frac{1}{r\sqrt{n}}) \\ 0 & \text{w.p. } 1 - \tilde{\Theta}(\frac{1}{r\sqrt{n}}) \end{cases}$$

The expectation of X can then be determined as follows.

$$\begin{aligned} E[X] &= E[\sum_{M\in U} \sum_{\{p \text{ destined for } M\}} X_M(p)] \\ &\leq E[\sum_{M\in U} \sum_{\{p \text{ destined for } M\}} \Theta(\log n) \times \tilde{X}_M(p)] \\ &= \sum_{M\in U} \Theta(\log n) \times \tilde{\Theta}(\frac{1}{r\sqrt{n}}) \times \lambda_M \leq \tilde{\Theta}(\sqrt{n}) \end{aligned}$$

Further, using a Chernoff bound, the high probability (i.e., $1 - 1/n^3$) result follows. We skip the details for brevity (we refer to [3], [16] for an analogous bound). ■

Each of the sets U_i satisfy the conditions of the above claim, and hence each produces a load less than $\tilde{\Theta}(\sqrt{n})$ with probability at least $1 - \frac{1}{n^3}$. Let the load from mobiles in U_i be X_i.

$$\begin{aligned} &\mathbb{P}\{\sum_{i=0}^{F} X_i \geq \tilde{\Theta}(\sqrt{n})\} \\ &\quad <_{(a)} \mathbb{P}\{\exists i : X_i \geq \tilde{\Theta}(\sqrt{n}/F)\} \\ &\quad = \mathbb{P}\{\exists i : X_i \geq \tilde{\Theta}(\sqrt{n})\} \\ &\quad <_{(b)} \sum_{i=0}^{F} \mathbb{P}\{X_i \geq \tilde{\Theta}(\sqrt{n})\} \\ &\quad \leq (F+1)n^{-3} \leq \tilde{\Theta}(n^{-3}). \end{aligned} \tag{4}$$

By the pigeonhole principle, if the sum of X_i is greater than $\tilde{\Theta}(\sqrt{n})$, one of the X_i must be greater than $\tilde{\Theta}(\sqrt{n}/F)$, giving inequality (a). The expression (b) follows from a union bound.

Finally, consider mobiles within radio range of T. Note that the total rate that mobiles can receive within this a radio range area is at most $\tilde{\Theta}(\sqrt{n})$, and packets for these mobiles will only enter the area once (since once they enter, they are received). Then clearly, the load here is at most $\tilde{\Theta}(\sqrt{n})$.

This establishes that the load on T is sufficiently low with high probability. The final step is to extend this and show that all tiles will have a low load with high probability. Let $X(T)$ denote the load at tile T. We have:

$$\begin{aligned} &\mathbb{P}\{\exists T : X(T) \geq \tilde{\Theta}(\sqrt{n})\} \\ &\leq \sum_T \mathbb{P}\{X(T) \geq \tilde{\Theta}(\sqrt{n})\} \\ &\leq \sum_T \tilde{\Theta}(n^{-3}) \\ &\leq \tilde{\Theta}(n \times n^{-3}) \\ &= \tilde{\Theta}(\frac{1}{n^2}). \end{aligned} \tag{5}$$

By Borel-Cantelli's Lemma, we have the load is feasible almost surely. ■

Thus, in the light of Theorems 3.1, 3.2 and 3.3, we have the following result.

Theorem 3.4: For any achievable traffic vector Λ (in the presence of non-ergodic mobility), the routing under MobileSearch is stable when a load of $\tilde{\Theta}(\Lambda)$ is imposed.

Proof: From Theorem 3.3, we see that for rates in the region $\tilde{\Theta}(\Lambda)$, the load on any tile of the network is no more than $\tilde{\Theta}(\sqrt{n})$ almost surely. This holds even for arbitrary movement of the mobiles because the cut-set bounds in Lemma 3.2 apply as if the mobiles were always static in the worst case positions. Since the network can support $\sqrt{n}$ packets per time-slot, this load is feasible. It can be shown there is a schedule that stabilizes the routing in this case.

In addition, by Theorem 3.2, we see that all packets either the reach their destination or the mobile is always contained in the 'active-grid' for each iteration of the algorithm. The routing is stable, so the algorithm will always continue until either the packet is received, or the approximately $(1/2)\log n$ iterations pass and the final active grid is reached. Since the final 'active-grid' is smaller than radio range and the mobile must be within it, the packet will get received in this case as well.

Therefore, we can conclude that the MobileSearch Algorithm will successfully route packets to their destination with a stable load imposed on the network. ■

IV. Discussion and Conclusion

In this paper, we showed that a mobile node can be reached by a path that is within a $\tilde{\Theta}(1)$ factor of the straight-line path to it from the source node. Our scheme utilized 'advertisements' by mobile nodes to aid the sequential search process - i.e., at each iteration the 'advertisement' packet helped guide the search into the correct sub-region. We also demonstrated that the sequentially finer grids did not overload any tile.

In our model, we assumed that there exists a unique source node A_i for each mobile destination M_i which is a $\Theta(1)$ distance away from $\mathcal{S}_i$. We note that the model can immediately be extended to multiple sources providing a sum rate λ_{M_i} to mobile M_i. This is due to the fact that independent grids are constructed for each packet, and the bounds on allowable distributions of source nodes are the same.

The restriction on the distance between the source and its corresponding mobile (i.e. the restriction that $d(\mathcal{S}_i, A_i) \sim \Theta(1)$) can be relaxed to allow all source-destination separations by modifying the algorithm as follows:

A. Modifications to MobileSearch

Instead of constructing large grids (of size $\Theta(1)$ in the very first iteration, the algorithm constructs a small grid of the size of the radio range and waits for the mobile's advertisement. If the mobile's location is outside the smaller grid, or the advertisement did not reach the current grid within a given timeout, the algorithm iterates to the next larger grid (of twice the perimeter). This proceeds with geometrically increasing grid size until a grid receives an advertisement from the mobile indicating that it is inside the current 'active-grid'. Once the mobile is trapped inside the active region, the algorithm proceeds as in Section III-A.

Notice that by this modification, the footprint of the grids creates no more than a $\tilde{\Theta}(1)$ factor larger load as compared to the optimal path (possibly multi-path) from the source to the mobile, i.e., the transport capacity imposed on the network is not increased significantly.

Acknowledgment: This work was partially supported by NSF Grants CNS-0347400, CNS-0519535, CNS-0626558, CNS-0721380 and the Darpa ITMANET program.

References

[1] 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.

[2] D. Braginsky and D. Estrin. Rumor routing algorithm for sensor networks, 2001. Available at http://lecs.cs.ucla.edu/˜estrin/.

[3] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah. Throughput delay trade-off in wireless networks. In *Proc. of IEEE Infocom*, March 2004.

[4] M. Grossgluauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In *Proceedings of IEEE Infocom 2001*, Anchorage, Alaska, April 2001.

[5] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. on Info. Theory*, IT-46(2):388–404, March 2000.

[6] 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.

[7] B. Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proc. of the ACM/IEEE International Conf. on Mobile Comp. and Networking*, page 243, August 2000.

[8] 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.

[9] Michael J. Neely and Rahul Urgaonkar. Cross-layer adaptive control for wireless mesh networks. *Ad Hoc Netw.*, 5(6):719–743, 2007.

[10] G. Pei, M. Gerla, and T.-W. Chen. Fisheye state routing: A routing scheme for ad hoc wireless networks. In *Proceedings of IEEE ICC*, New Orleans, LA, June 2000.

[11] 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.

[12] 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.

[13] S N. Diggavi, M. Grossglauser, and D. N. C. Tse. Even one-dimensional mobility increases capacity of wireless adhoc networks. In *IEEE International Symposium on Information Theory (ISIT), Lausanne*, page 352, 2002. inproceedings.

[14] Devavrat Shah and Sanjay Shakkottai. Oblivious routing with mobile fusion centers over a sensor network. In *INFOCOM*, pages 1541–1549, 2007.

[15] G. Sharma, R. Mazumdar, and N. Shroff. Delay-capacity tradeoffs in mobile ad-hoc networks: A global perspective. In *Proc. of IEEE Infocom*, 2006.

[16] S. Subramanian and S. Shakkottai. Geographic Routing with Limited Information in Sensor Networks. In *Proceedings of Information Processing in Sensor Networks*, pages 269–276, April 2005.

[17] S. Subramanian, S. Shakkottai, and P. Gupta. On optimal geographic routing for networks with holes and non-uniform traffic. In *Proceedings of IEEE Infocom*, 2007.

[18] S. Subramanian, S. Shakkottai, and P. Gupta. Optimal geographic routing for wireless networks with near-arbitrary holes and traffic. Bell Labs Technical Report, Jul 2007. Available at http://www.ece.utexas.edu/∼ssubrama/info08long.pdf.

[19] Lei Ying and R. Srikant. Optimal delay-throughput trade-offs in mobile ad-hoc networks: Hybrid random walk and one-dimensional mobility models. *CoRR*, abs/0705.0326, 2007.