Link-diversity Routing: A Robust Routing Paradigm for Mobile Ad Hoc Networks

TIK Report 268

Vincent Lenders*, Rainer Baumann†
*Department of Electrical Engineering
Princeton University, NJ, USA
†Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland

***Abstract*— We present *link-diversity routing*, a routing paradigm that increases the path resilience in mobile ad hoc networks. Link-diversity routing chooses each hop of a packet's route, so that the choice reflects the amount of outgoing links at the intermediate hops. This choice maximizes the opportunities to make progress at every hop in the presence of link failures. As a result, link diversity routing takes paths which are less prone to fail due to individual link failures than traditional routing. We develop a loop-free and distributed link-diversity routing algorithm. The algorithm is based on an analogy from the heat theory which consists of routing packets along the steepest gradient of a temperature field that is created by the destination nodes. We perform simulations of our algorithm with a DSDV-based implementation. Our simulations of 802.11b networks with up to 10000 nodes which are moving according to a mobility model we inferred from geographical data of a large existing city show that link-diversity routing increases considerably the path resilience and also the overall performance. For nodes moving at car speeds in our simulations, link-diversity routing increases the end-to-end packet delivery ratio compared to traditional routing by a factor of up to four.**

I. INTRODUCTION

Mobile wireless multi-hop networks typically use *best effort* routing techniques. That is, each router in the network maintains a forwarding table that includes a set of possible links to reach a specific destination[1]. When a packet has to be forwarded, a node looks up in its forwarding table the preferred link according to the network routing metric, and forwards the packet over this link. If for any reason this link is currently unavailable (for example, the node is temporarily down or has just moved away before the change has been captured and updated by the routing protocol), an alternative link from the forwarding table is selected and used to forward the packet. When there are no alternative links in the forwarding table, the packet is simply dropped and the loss is assumed to be handled by the upper layers.

This model is in itself very robust to link failures as it works even in the presence of link failures. However, the probability that a packet eventually arrives successfully at the destination strongly depends on the number of outgoing links at each hop along a path to the destination. This issue is best illustrated in Figure 1. In the top network topology, a source node is connected to a destination node via a single chain of intermediate nodes. If any link on this chain becomes unavailable, a packet from the source to the destination node cannot be delivered. In the bottom of the figure, all nodes have two outgoing links in their forwarding table over which they can reach the destination. In contrast to the previous case, a packet from the source node can now be delivered successfully in the presence of individual link failures.

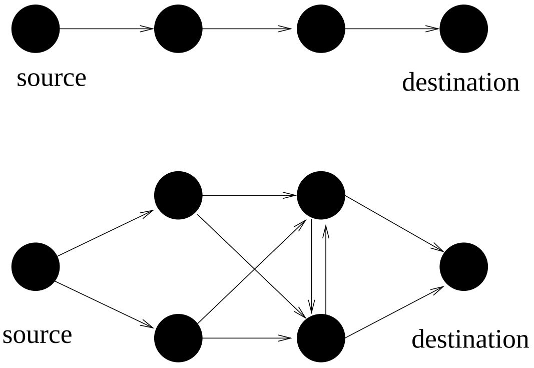

Fig. 1. The top network topology is more sensitive to link failures than the bottom one because each node has only one link to reach the destination. If only one link fails, the source can no longer reach the destination. In the bottom topology, the source node can still reach the destination when particular links are unavailable.

Existing routing protocols are agnostic to the link diversity and typically optimize for a metric like the number of hops [26], [18], [27], [8], [13], the link qualities [28], [10], [1], or the throughput [9]. These approaches are effective in rather static networks where the link failure probability is low. However, as illustrated before, in dynamic and possibly mobile networks, ignoring the alternative link opportunities along a path often leads to dead-ends.

In this paper, we propose *link-diversity* routing, a novel routing paradigm that increases the path resilience in the presence of frequent link failures. This is achieved by choosing each hop of a packet's route such that each hop has a high number of opportunities to forward the packet. As a consequence, it is possible that packets will not be delivered

---

[1] A destination can be a single node or a set of nodes aggregated as a whole network.

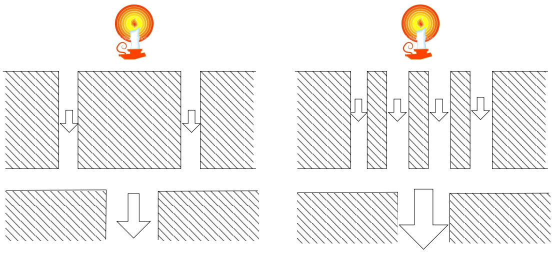

Fig. 2. The heat flow at the bottom of the left arrangement is lower than at the bottom of the right arrangement because of the difference in path connectivity.

over the shortest path to the destination.

To implement link-diversity routing, we present the *finite difference method routing* (FDMR) algorithm which is based on an analogy from the heat theory. Our algorithm exploits the fact that the heat flow generated by a heat source depends on the physical interconnection of the propagation media (see Figure 2). In our routing context, this means that highly interconnected network regions, or nodes with a high-link diversity, will allow the heat to dissipate better than sparsely connected regions, or nodes with a low link diversity. As a result, our forwarding scheme which consists of forwarding packets towards the ”warmest” path, routes packets along paths which exhibit a high link-diversity.

The main contributions of this paper are the following:

- We propose the idea of link-diversity routing to improve the path resilience in networks with frequent link failures.
- We present the design of a routing model inspired from heat theory that exploits link-diversity. Additionally, we provide the FDMR algorithm, a loop-free routing algorithm that relies on the finite different method to distribute heat information in the network in a distributed way.
- We compare the performance of our DSDV [27]-based FDMR implementation with the original DSDV protocol that uses a the minimum hop-count as routing metric. Our simulations show that link-diversity routing manages to increase the end-to-end packet delivery probability in mobile scenarios by a factor of up to four. This improvement is achieved without considerably increasing the delivery path length.

The rest of this paper is organized as follows. In the next section, we present the heat-based routing model and the FDMR algorithm. In Section III, we describe our implementation of the FDMR algorithm which relies on the DSDV routing protocol. Section IV presents our simulation results. We review related work in Section V and conclude in Section VI.

II. MODEL FOR LINK-DIVERSITY ROUTING

In this section, we present our heat-inspired model for link-diversity routing. We first describe the key ideas of the heat-based routing paradigm. Then, we introduce the FDMR routing algorithm which we propose to establish temperature fields in a completely distributed manner and forward packets according to those fields. We also prove the key routing properties such as convergence and loop-freeness.

A. Heat-based Routing Paradigm

The key idea of link-diversity routing is to chose a packet's route so that the amount of outgoing links at each hop is sufficiently high to provide enough forwarding opportunities. To find such routes, we rely on a thermodynamic analogy. Consider the heat flow away from an ideal heat source. The flow intensity depends on how well a region is physically interconnected with the source (see Figure 2). To exploit this property for routing, we model the destination in a network as a heat source. Then, we evaluate the heat flow resulting from this heat source at all the nodes according to the network connectivity. For this, we determine the temperature of every node. The temperature distribution around a heat source is defined by the heat equation. The heat equation is a partial differential equation which in steady state follows the Laplace equation:

$$\Delta\phi = 0, \tag{1}$$

where Δ is the Laplace operator defined as $\nabla^2 = \nabla \cdot \nabla = \text{div grad}$[2], and ϕ is the temperature distribution. This equation basically means that the temperature distribution is a twice differentiable function. It also means that the field is a monotonously decreasing function.

As such, the heat equation defines an infinite number of possible solutions. An exact solution of the equation is one that satisfies the a priori known *boundary conditions*. We define the boundary conditions intuitively as follows. We set the temperature of the destination node x_d to $\phi(x_d) = 1$, corresponding to the constant temperature of an ideal heat source and the temperature value of the source node x_s to $\phi(x_s) = 0$, corresponding to the absolute lowest possible temperature. These two boundary conditions together with the heat equation now define a unique temperature field distribution which has a maximum value at the destination node and is minimal at the source node. Furthermore, it can be shown that since the Laplace equation is a monotonously decreasing function, the distribution has no local maxima. This is a necessary condition for the routing to work as we will see later.

Figure 3 shows the temperature field for an example network. The temperature of the nodes are given by the numbers inside the circles representing the nodes. The temperature of the boundary conditions are set to 1 (for the destination) and 0 (for the source).

Finding a path from the source to the destination is a gradient search problem. Since a field distribution following the heat equation can never have a local maximum, any ascending gradient from the source (represented with a thin arrow in Figure 3) is a valid path towards the destination. Therefore, loop-free packet routing can be implemented in a hop-by-hop way by forwarding a packet at every node to any neighbor which has a higher temperature value. In our approach however, unless a link along the steepest gradient

---

[2] where div is the divergence and grad is the gradient.

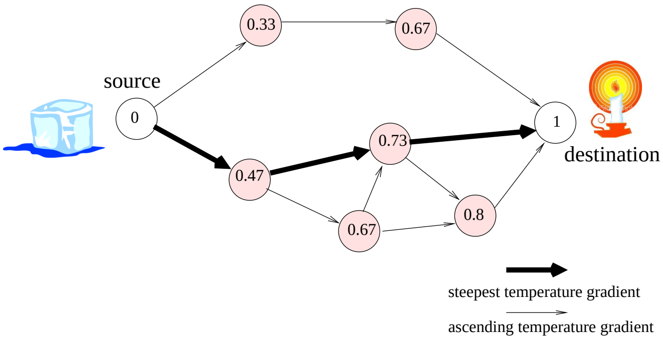

Fig. 3. Example temperature field. The steepest gradient from the source to the destination is along the path with the highest link diversity.

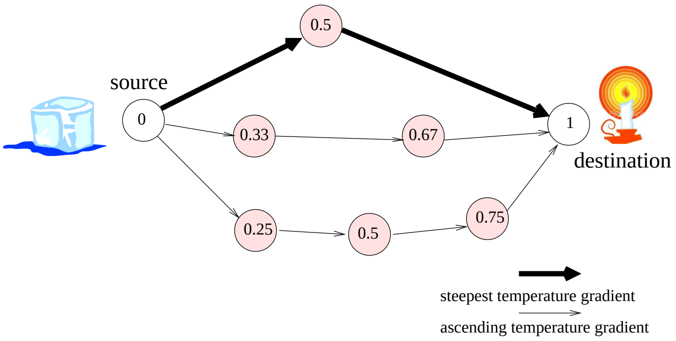

Fig. 4. Example temperature field: All paths have the same link diversity and the steepest gradient is along the shortest path.

is broken, we always forward packets along the steepest ascending gradient (represented with a thick arrow in the figure), corresponding to the neighbor with highest temperature. This way, the path a packet follows is analogous to the minimum-energy diffusion path of a particle in a real temperature field. Further notice that the chosen path towards the bottom is more resilient to link failures than the upper path.

Another interesting property of the routing model it that the path with the minimum hop-count is chosen if the link-diversity of different paths is the same. This is best illustrated in Figure 4, where there are three paths from the source to the destination all having the same link diversity but a different number of hops. As we can see the steepest gradient is along the path with the minimum number of hops.

B. The FDMR Routing Algorithm

Establishing routes with this model results in establishing a temperature field on the network nodes which solves Equation (1). To establish a temperature field, we require an algorithm which is by design scalable and robust. These two properties imply the following two requirements:

- The algorithm should be completely distributed.
- Every node should be able to calculate its own temperature locally, based on the temperature of *only* its direct neighbors.

We propose the FDMR algorithm which fulfills both requirements at the same time. The FDMR algorithm relies on

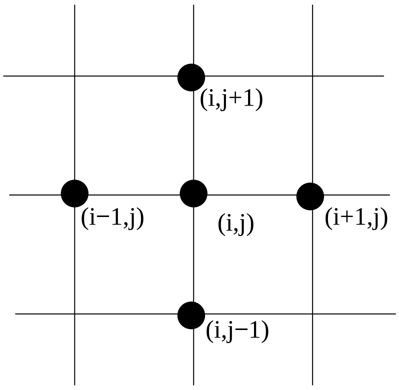

Fig. 5. The finite difference method on a grid: $\phi_{i,j} = \frac{1}{4}(\phi_{i+1,j} + \phi_{i-1,j} + \phi_{i,j+1} + \phi_{i,j-1})$

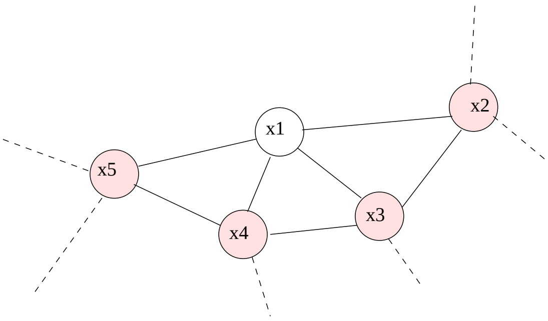

Fig. 6. The FDMR algorithm: Nodes calculate their temperature as the average temperature of their neighbors: $\phi_{t+1}(x_1) = \frac{1}{4}\ \phi_t(x_2) + \phi_t(x_3) + \phi_t(x_4) + \phi_t(x_5)$

a numerical technique called the *finite difference method* [22] to evaluate the solution of partial differential equations on a grid. The finite difference method is an iterative technique. As a basic principle, the method relies on the fact that solutions to Laplace's equation are harmonic functions and thus satisfy the mean value theorem of potential theory. According to this theorem, the value at a point is equal to the arithmetic average of its values on a boundary surrounding this point. For example, in a 2-dimensional square grid, the value at a point is equal to the average values of the four immediately neighboring points in the grid (see Figure 5). The way the finite difference method works is by iteratively approximating the value of the points on the grid as the arithmetic average of the neighboring points until the values have converged to the final solution. It can be shown that this method always converges to Equation (1) in a bounded number of iterations.

Our application of this method to evaluate the temperature distribution on a network of nodes is as follows. We define the grid as given by the network topology. That is, every node in the network is a point in the grid and every link in the network corresponds to an edge in the grid. Note that for an arbitrary network topology, the grid might not be a regular square grid but this is not a requirement for the finite difference method to work.

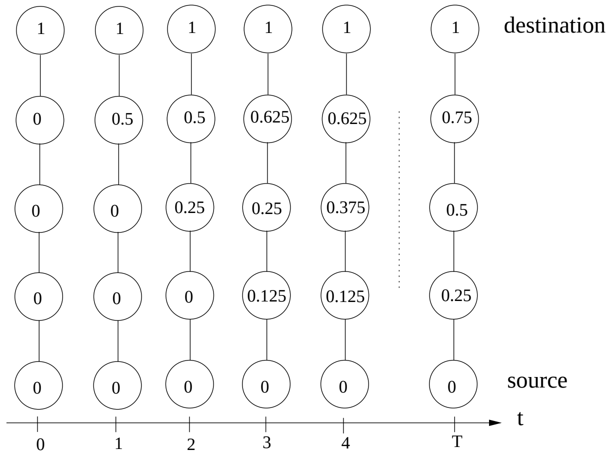

Fig. 7. Basic operation of the FDMR algorithm. At iteration step T, the algorithm has converged and the temperature distribution follows Equation (1).

We define the FDMR algorithm as follows. Let $X = \{x_1, \cdots, x_n\}$ be the set of nodes in the network and denote $\{x_k; k \in nbr(x_i)\}$ as the set of nodes which are neighbors of x_i (i.e., there exists a link between x_i and x_k). All nodes except the source and the destination nodes (which define the boundary conditions) start with an initial temperature value $\phi_{t=0}(x_i) = 0; x_i \in X \backslash \{x_s, x_d\}$ and calculate their own temperature at iteration step $t+1$ as the average of the temperature of their direct neighbors at iteration step t (see Figure 6 for an example):

$$\phi_{t+1}(x_i) = \begin{cases} \frac{\sum_{k \in nbr(x_i)} \phi_t(x_k)}{|nbr(x_i)|} & , |nbr(x_i)| > 0 \\ 0 & , |nbr(x_i)| = 0 \end{cases} \quad (2)$$

The temperature of the source node x_s and destination node x_d are the boundary conditions and thus constant over all iterations:

$$\phi_t(x_s) = 0; \forall t \geq 0 \quad (3)$$

$$\phi_t(x_d) = 1; \forall t \geq 0 \quad (4)$$

The operation of the algorithm is shown for a simple example in Figure 7. At iteration step $t = 0$, all nodes expect the destination have a temperature value of 0. At each next iteration step, the nodes recalculate their temperature according to the new temperature of their neighbors. At iteration step T, all values have converged.

C. Properties of the Routing Algorithm

We now prove three fundamental routing properties of the FDMR algorithm.

Theorem 1: The temperature field constructed by the FDMR algorithm converges to Equation (1) in a bounded number of iterations:

$$\Delta\phi_t(x) = 0, t < \infty. \quad (5)$$

Sketch of Proof: A detailed proof of this theorem is outside the scope of this paper. However, since our algorithm relies on the finite difference method to calculate the temperature field, the theorem holds if we prove that the finite difference method works on any irregular grid. This has been shown in [24]. ■

Theorem 2: In a converged temperature field constructed by the FDMR algorithm, a packet from the source forwarded along any ascending gradient will eventually reach the destination.

Sketch of Proof: To guarantee that any gradient from the source leads to the destination, two conditions must be met: (i) the temperature of the destination must be the maximum value in the field and (ii) the temperature may not have any local maxima. Condition (i) is always true since for a node to have temperature value greater than one, it must have at least one neighbor with a value greater than one. However, since at iteration step $t = 0$ there are no nodes with a value greater than one, it is not possible that at any iteration step $t \geq 0$ it occurs. Condition (ii) holds because any solution of the Laplace equation is an harmonic function. An harmonic function can never have a local maximum or minimum.■

Theorem 3: Packet forwarding with the FDMR algorithm is loop-free.

Proof: We prove this with a counter example. Assume that a packet has traveled in a loop. Therefore, the packet must have traveled across a node twice. However, the FDMR algorithm requires forwarding to be along an ascending gradient. Therefore, the potential value at every hop must be strictly larger than at the previous hop. Hence, it is not possible that a packet travels more than once through a specific node.■

III. IMPLEMENTATION OF THE FDMR ALGORITHM

In this section, we describe our implementation of the FDMR routing algorithm. Since our algorithm operates with iterative updates between the neighbors, it is a natural choice to implement it as a distance vector [3] routing protocol. For our implementation, we used the DSDV [27] routing protocol, a protocol designed for wireless ad hoc networks. In the following, we describe how we adapted the DSDV routing protocol and how packet forwarding is done.

A. Establishing Temperature Fields

Nodes running a distance vector protocol periodically exchange routing messages with their neighbors. These periodic messages include a cost per known destination. These costs represent the network ”distance” of each node to reach the destination. Looking at the cost of the direct neighbors, each node then estimates its own cost through a well defined method identical for all nodes. In the traditional DSDV implementation, the destination sets its own cost to zero and each node in the network calculates its own cost by adding one to the cost of the neighbor with the lowest cost. Hence, after convergence, the cost of each node is the distance in hops to the destination node.

In our implementation of the FDMR algorithm, the basic mechanism remains the same but the calculation of the costs is different compared to the original protocol. In our implementation, the destination node sets its cost to 1 and the source node to 0. All other nodes calculate their own cost as the average cost of their neighbors (see Equation 2). Hence, after convergence, the cost of each node corresponds to the temperature of the destination field which is a value between 0 and 1.

B. Packet Forwarding

The basic forwarding mechanism of the distance vector routing algorithm consists of decreasing the cost at each hop a packet is forwarded. This way, packets will eventually arrive at the destination which has the lowest cost. In the traditional minimum hop-count approach, each node sends the packet to the neighbor having the lowest cost. However since in our approach, the cost definition is inverted, forwarding is along increasing temperatures. In particular, we always forward packet along the steepest ascending gradient corresponding to the neighbor with the highest temperature.

When a link breaks, a certain amount of time is required for the routing protocol to react to the topology change and update the routing state in the network. To avoid that all paths using this particular link are temporarily down during this time, it is therefore necessary to react to such events at the forwarding level. Therefore, when a link from a used path is temporarily unavailable, the intermediate node which detects the failure (link failures can typically be detected when no periodic updates have been received from the neighbor within a certain period of time or in the presence of a missing link-layer acknowledgment after some amount of link-layer retransmissions) looks up in its local cache for an alternative neighbor with a higher temperature. If there is any neighbor with a larger temperature, the traffic is now forwarded to the neighbor with the highest temperature among them. Otherwise, the packets for this destination cannot further proceed towards the destination and should be temporarily cached until the protocol has managed to update the topology change, or dropped if the cache is full before an update has occurred.

IV. EVALUATION

This section examines the performance of the FDMR routing algorithm.

A. Experimental Methodology

We assess the benefit of link-diversity routing by comparing the performance of the FDMR algorithm with the traditional DSDV routing protocol that uses the minimum hop-count metric. Since our implementation of the FDMR algorithm relies on the DSDV protocol to establish the temperature fields, it allows for a fair comparison where the control message overhead is identical for both protocols.

The only difference between the FDMR-enhanced and the traditional DSDV protocol lays in the way how the costs are calculated based on the periodic update information received from their direct neighbors. In the FDMR approach, each node uses the arithmetic average of the cost of its neighbors, whereas in the minimum hop-count approach, each node calculates its cost as one hop plus the hop-count of the closest neighbor to the destination.

In a first step, we present results from an idealized simulation study using random graphs. These simulations are not meant to provide quantitative measures on system-level performance metrics but rather more profound understanding on the behavior of both algorithms. Then, in a more realistic simulation study using Glomosim [32], a wireless network simulator incorporating sophisticated models for the physical and MAC communication layers, we analyze the performance of both schemes for concrete scenarios with mobility.

B. Simulations with Random Graphs

We now present our simulation results with random graphs. The considered graphs are random unit disk graphs generated by assigning the position of vertices randomly over a two-dimensional square area. There exists an edge between two vertices in the graphs if the geometric distance of the vertices is smaller than a fixed threshold value. This model is an ideal simplification for the network topology that arises from randomly placed nodes sending at the same transmission power in a free space environment.

1) Robustness: To capture the benefit of link-diversity routing in terms of robustness, we look at the ability of both protocols to redirect traffic around broken links during the moment directly after a topological change occurs and the moment until the protocol had time to react to the topological changes and update its state information. We do this by measuring the packet delivery ratio in simulations where we allow both protocols to convergence to their final state but define random edges in the graphs as `failed`, meaning that no packet can be forwarded over those particular links. In wireless networks, such link failures could for example arise from short-term fading effects resulting in temporary unavailable links. Or in mobile networks, some nodes might move causing individual links to break when two nodes move outside communication range before the routing protocol was able capture the failure and update it. Using this approach, we are able to see how well each algorithm manages to handle broken or temporarily unavailable links at intermediate hop when links along the least cost paths are unavailable.

We plot a representative snapshot of our simulation results in Figure 8 and 9. These particular simulations were obtained from random graphs generated with a square side length of 2500 meters and a threshold range of 250 meters. As communication model, we chose randomly selected vertices acting as source nodes and sending packets to randomly selected vertices acting as destination nodes. Figure 8 shows the result using a graph size of 400 vertices. The x-axis represents the ratio of edges which we defined as `failed` after convergence of both algorithms. A value of 0.1 means that 10 percent of the total edges are `failed`. As we can see, the FDMR algorithm achieves a much higher packet delivery probability than the minimum hop-count based algorithm as the ratio of `failed` edges increases. To capture the effect of the network density on this metric, Figure 9 shows the result for the same simulations

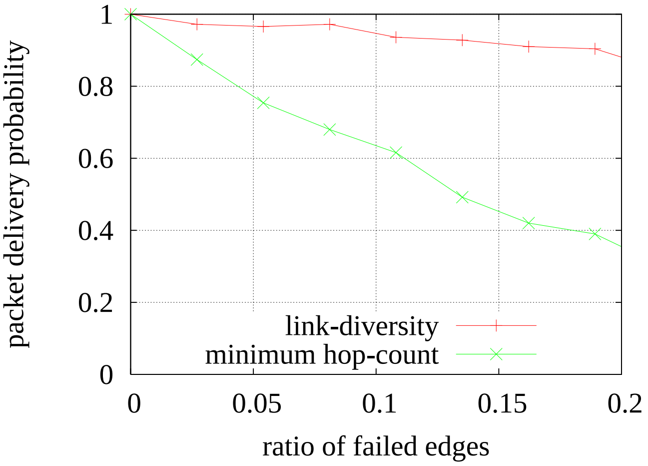

Fig. 8. Robustness: Packet delivery probability versus the fraction of failed edges in a graph of 400 vertices.

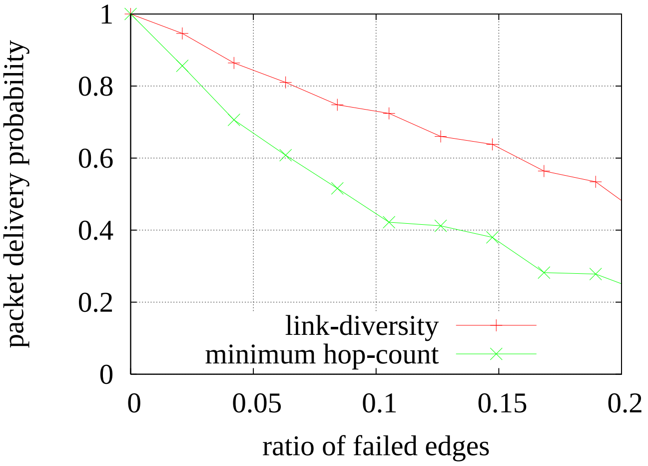

Fig. 9. Robustness: Packet delivery probability versus the fraction of failed edges in a graph of 200 vertices.

but with a graph size of 200 nodes instead of 400 nodes. With a smaller graph size, both algorithms achieve a lower delivery probability. However, the difference between the two curves is smaller than before as the ratio of `failed` edges increases. This is due to the fact that a less dense graph has fewer edges and inherently a smaller link diversity towards the destination. Therefore, we can expect the most benefits of link-diversity routing in rather dense networks.

2) *Convergence:* Another important metric is the convergence time of the algorithms. So far, our simulations did not account for this metric since we performed them assuming that the algorithms had totally converged. In a first attempt to quantify the algorithms' convergence, we define the convergence time as the number of periodic update intervals required from scratch to achieve a routing state which no longer changes at any network nodes as further updates are sent. Since with the FDMR algorithm, routing is not affected by the absolute node temperature values but rather the differences (or gradients) between those values, we measure convergence of this protocol by evaluating the gradient at all the links and determining when the direction of these gradients no longer changes.

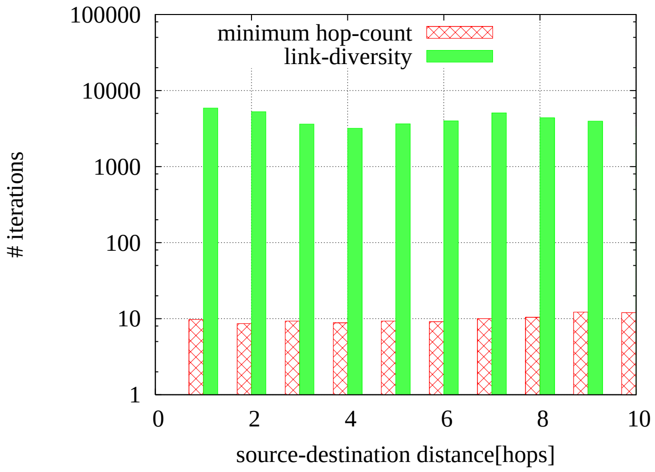

Fig. 10. Convergence: number of iterations required until the routing state has converged at every node.

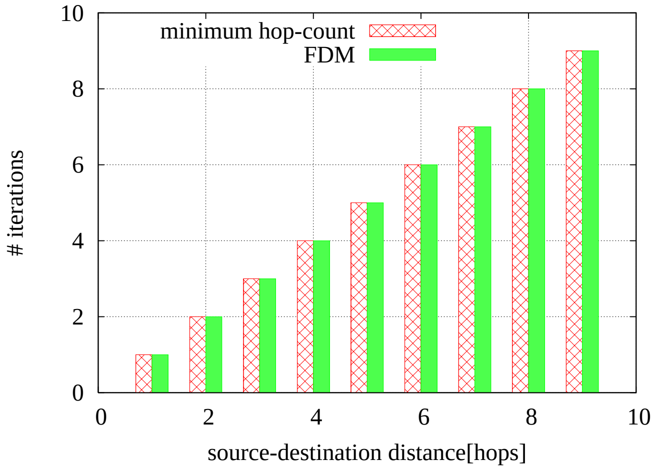

Fig. 11. Convergence: number of iterations required until a path is established between a source and a destination versus the distance between the source and the destination.

With this definition, the convergence time is given in Figure 10. On the vertical axis, we plot the number of iterations defined as the number of required periodic update intervals. The horizontal axis represents the path length between a source node and a destination node. Note that the results are from simulations using the same parameters as previously and a graph size of 400 nodes. At a first glance, this result appears quite discouraging since the convergence time of the FDMR algorithm is almost three power orders higher than the minimum hop-count algorithm. However, recall that this convergence definition requires that the routing state no longer changes at any node in the graph. If alternatively, we consider as convergence metric the number of update intervals required for a route to be available between a source and a destination

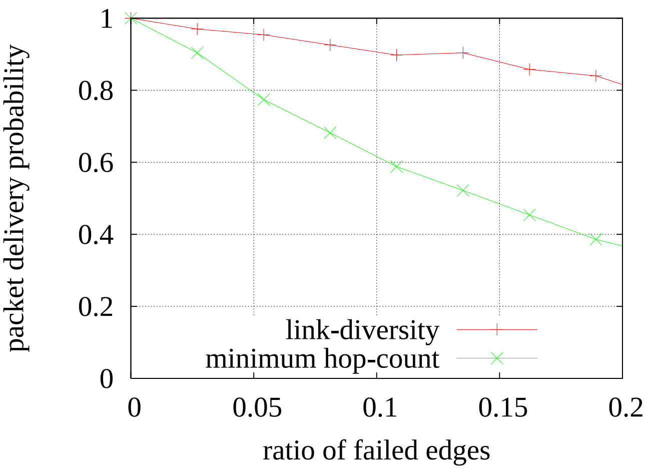

Fig. 12. Robustness after limited convergence (400 vertices). Link-diversity routing still achieves a higher packet delivery probability than the minimum hop-count approach.

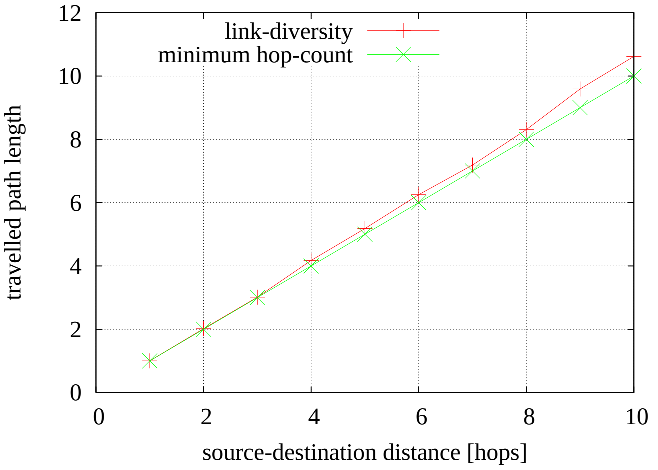

Fig. 13. Path stretch: Traveled path length of data packets versus the distance of the sending node to the destination.

node, the picture changes completely as can be seen in Figure 11. In this figure, we show the number of required update intervals (on the y-axis) versus the source-destination distance (on the x-axis) until a path is established between the source and the destination. With this different convergence definition, both algorithms converge at a time which is directly proportional to the source-destination distance. Therefore, the FDMR algorithm converges across the direct path to the destination as quickly as the minimum hop-count algorithm and only requires longer to converge at the nodes which are away from it.

Since the FDMR algorithm requires a very large amount of iterations to completely converge but manages to find a path to the destination in the same linear amount of time as the minimum-hop count algorithm, the remaining question is whether the FDMR algorithm still achieves higher delivery probabilities for a number of iterations in the order of the source-destination distance. To answer this question, we repeated the robustness simulations presented earlier but did not use a converged routing state. Rather, we limited the number of iterations for both algorithms to a value which was identical to the source-destination distance. The result for a graph size of 400 nodes is plotted in Figure 12. The delivery probability for the minimum hop-count algorithm remains almost the same and decreases slightly faster for the FDMR algorithm than in Figure 8. Still, there is a significant difference between the FDMR and the minimum hop-count algorithm. This result shows that even with a small number of iterations, the FDMR routing algorithm manages to exploit well the link-diversity and increase the routing robustness. A slightly simplified explanation for this result is that there are multiple link-disjoint but not necessarily node disjoint short paths between the source and the destination and that these converge equally fast.

3) Path Stretch: Finally, it is also worth considering the path stretch induced by link-diversity routing. It is for example not desirable to increase the packet delivery probability if it comes at the cost of extremely long paths and thus also high delivery delays. In Figure 13, we plot the traveled path length of data packets (on the vertical axis) versus the shortest source-destination distance (on the horizontal axis) for the same scenario as considered in Figure 8 but without any `failed` links. The minimum hop-count algorithm represents a linear curve as it always delivers packets over the shortest path. On the other hand, the paths chosen by the FDMR algorithm are only marginally longer than the shortest distance. Thus, the FDMR algorithm significantly increases the packet delivery probability without significantly increasing the average traveled path length.

C. Simulations of Mobile Networks

In the following, we study the performance of the FDMR algorithm in a more realistic context using the Glomosim simulator [32]. We consider the performance in a city-wide ad hoc network in which the nodes are mobile and communicating over WLAN. These types of networks could be used for many different purposes in the area of urban inter-vehicle communication or person-to-person communication. Our simulation setup and assumptions are described next.

1) Simulation Setup: We used a IEEE 802.11b network with a capacity of 11 Mbps and a nominal wireless range of 250 meters. As MAC protocol, we used the 802.11 DFWMAC-DCF w/RTS/CTS and as propagation model two-ray ground. Due to the large network sizes we use, we were unable to simulate the effect of intermediate buildings.

We used two different mobility models: the steady-state random trip mobility model [15] on a network of streets and the random waypoint mobility model [17][3]. In both models, the nodes move with constant speeds and without pausing on a square of side length 10km by 10km. However, in the random trip mobility model, nodes move on vectorized maps

---

[3] The random waypoint mobility model has shown to have non-desired behavior [4] when not well parameterized. We follow the guidelines as proposed in [23] to avoid such effects.

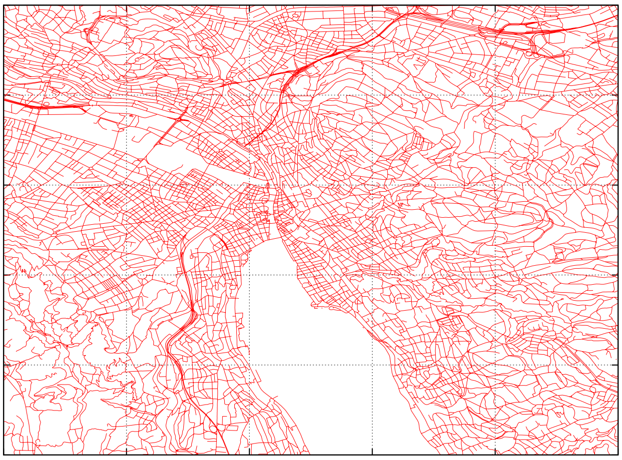

Fig. 14. Vectorized street map of a city (10km by 10km).

which we extracted from a geographic information system (GIS). The vectorized map of a city center for which we present our results in this paper is shown in Fig. 14. In the random waypoint mobility model, the nodes move from randomly chosen waypoint in the square to another waypoint on a straight line. The random waypoint model is far less realistic but included as a reference because it is often used in the literature. We differentiate two concrete scenarios. In the pedestrian scenario, nodes move with a speed uniformly distributed in the range $[1-4]m/s$. In the car scenario, nodes move with speeds in the range $[10-20]$ m/s. These speeds corresponds to typical pedestrian and car speeds in a typical city.

As traffic model, we send constant bit rate traffic from randomly chosen source nodes to randomly chosen destination nodes. All packets are 1024 bytes long. All simulations have a duration of at least 10000 seconds and are always an average over at least 20 runs with different random seeds.

2) Results: The performance results for nodes moving along the streets of a large city are shown in Figure 15. In the upper figure, 10000 nodes are used resulting in an average node degree of approximately 11. In the lower figure, 5000 nodes are used, leading to an average node degree of 5.5. In both settings, 500 nodes are active traffic sources out of which approximately two third are sending simultaneously on average over the simulation time. As a reference, we also plot the packet delivery ratio when the nodes are not moving (static scenario). As we can see, the ratio in the static scenario is close to 100% which means that most packet losses occur in the mobile scenarios because of the mobility of the nodes and not due to other effects like congestion or interference.

We conclude that the packet delivery ratio is clearly better with FDMR compared to minimum hop-count routing. Another interesting observation is that the performance with minimum hop-count does not significantly gets better as the node node density increases. This is different for the FDMR algorithm. With our algorithm, the performance gets better as the average node degree increases. This is because the FDMR algorithm exploits the link-diversity which becomes larger when the node degree increases.

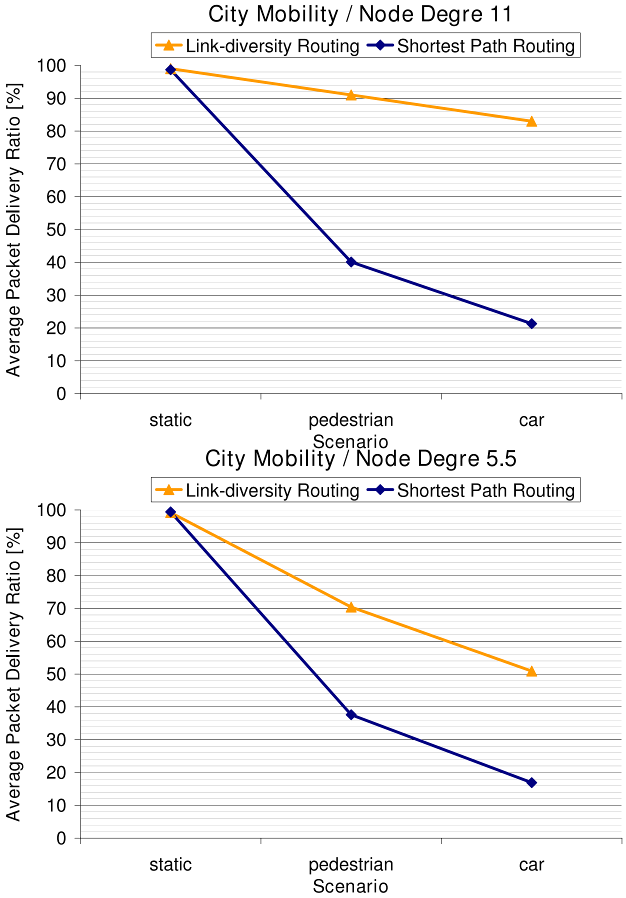

Fig. 15. Street mobility model: Performance comparison of DSDV with traditional minimum hop-count and with link diversity.

The simulation results from the random waypoint model are given in Figure 16. The trends we observe with this model are the same as with the previous mobility model. Overall, the packet delivery ratio is slightly better with this model. One of the main reason is that the node distribution is not uniform. Nodes tend to concentrate around the center of the square [4]. As a result, the average path length is smaller than with the previous mobility model, and on average, paths break less frequently.

V. Related Work

Traditional routing protocols for mobile ad hoc networks typically optimize strictly for a metric like the number of hops [26], [18], [27], [8], [13], the link quality [28], [10], [1], or the expected throughput [9].

The authors of [5], [19], [16], [7] exploit the link diversity in multi-hop routing to cope with lossy links and fading effects. These proposed techniques integrate MAC-layer and routing techniques by broadcasting (or anycasting) packets to all nodes in transmission range and letting those decide whether to further forward the packet or not. These approaches are complementary to our routing algorithm since they could profit from routing over paths with a high link diversity and vice versa.

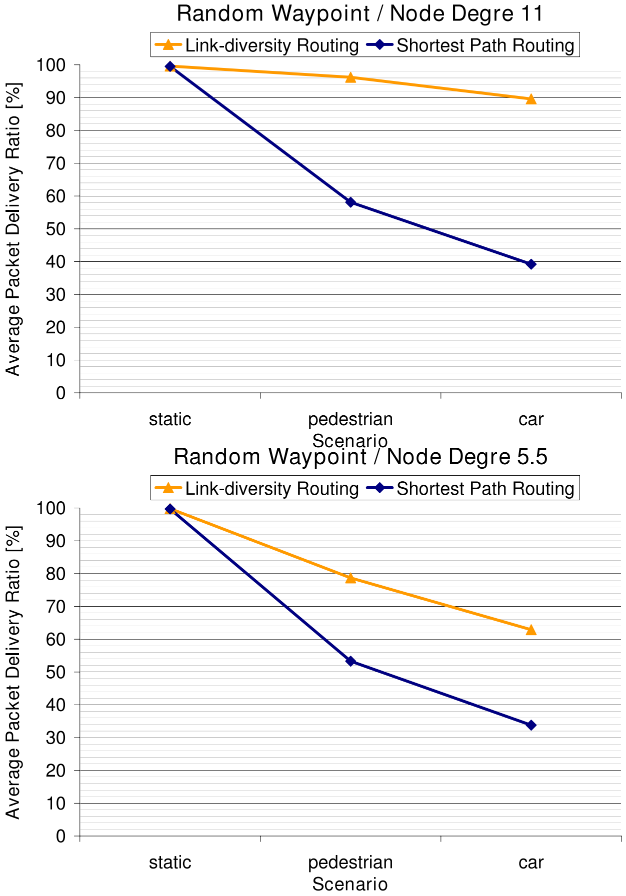

Fig. 16. Random waypoint mobility model: Performance comparison of DSDV with traditional minimum hop-count and with link diversity.

Braided multipath routing [12] identifies multiple routes, using one as a primary and switching if the primary fails. Opportunistic multipath scheduling [6] splits traffic over multiple paths, adaptively favoring paths that provide low delays. Tsirigoris and Haas [30] propose to use erasure coded fragments of each packet over disjoint paths in a mobile ad hoc network, in order to tolerate loss of some fragments due to fading or node movement. Link-diversity routing also exploits multiple paths, but selects them according to the amount of forwarding opportunities at each hop and must not ensure that the paths are disjoint.

The steepest gradient search method has been well studied in the past. This method has been extensively used for optimization problems and has had applications in diverse disciplines as routing in ad hoc networks [25], [20], load balancing in the Internet [2], data collection in sensor networks [14], [11], sensor node placement [29], guided navigation [21], or service discovery [31]. The basic forwarding principle of our approach which consists of forwarding along the steepest gradient is similar to these works. However, our distributed method to establish and maintain a potential field which is based on the finite different method is unique and uses only local information from the direct neighbors by mimicking how heat flow dissipates. Furthermore, to the best of our knowledge, we are the first to exploit the properties of fields in terms of link diversity.

VI. CONCLUSIONS

This paper presents link-diversity routing, a robust routing paradigm for mobile ad hoc networks. Link diversity routing increases the path resilience by choosing a packet's route so that the amount of outgoing links at each hop is maximized. Therefore, the selected hops have often multiple forwarding opportunities and can better cope with individual link failures caused by mobility and fading than traditional routing schemes.

We show that by modelling the destination as a heat source and routing along the steepest gradient of the temperature field created by this source, link-diversity routing can be implemented in a distributed and loop-free manner. We provide the FDMR algorithm which enable to calculate the temperature of a node in the network based only on the temperature of its direct neighbors.

Using simulations with static random graphs and Glomosim-based network simulations with node mobility, we show that link-diversity routing increases the end-to-end packet delivery ratio compared to traditional minimum hop-count routing by a factor of up to four (when nodes are moving at car speeds in a city). This improvement is achieved without significantly increasing the average path length and with the same control message overhead.

REFERENCES

[1] Sulabh Agarwal, Ashish Ahuja, Jatinder Pal Singh, and Rajeev Shorey. Route-lifetime assessment based routing (RABR) protocol for mobile ad-hoc networks. In *ICC (3)*, pages 1697–1701, 2000.

[2] Anindya Basu, Alvin Lin, and Sharad Ramanathan. Routing Using Potentials: A Dynamic Traffic-Aware Routing Algorithm. In *Proceedings of the ACM annual conference of the Special Interest Group on Data Communication (SIGCOMM'03)*, Karlsruhe, Germany, August 2003.

[3] Dimitri Bertsekas and Robert Gallager. *Data Networks*. Prentice-Hall, Englewood Cliffs, New Jersey, USA, 1992.

[4] Christian Bettstetter, Giovanni Resta, and Paolo Santi. The Node Distribution of the Random Waypoint Mobility Model for Wireless Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 2(3):257–269, 2003.

[5] Sanjit Biswas and Robert Morris. Exor: Opportunistic multi-hop routing for wireless networks. In *SIGCOMM*, 2005.

[6] C. Cetinkaya and E. Knightly. Opportunistic Traffic Scheduling over Multiple Network Paths. In *Proceedings of IEEE INFOCOM*, Hong Kong, September 2004.

[7] R. Roy Choudhury and N. Vaidya. MAC Layer Anycasting in Wireless Networks. In *Proceedings of the HotNets II Workshop*, Cambridge, MA, USA, November 2003.

[8] Thomas Clausen and Philippe Jacquet. Optimized Link State Routing Protocol. IETF Internet Draft, draft-ietf-manet-olsr-11.txt, July 2003.

[9] Douglas S. J. De Couto, Daniel Aguayo, John Bicket, and Robert Morris. A high-throughput path metric for multi-hop wireless routing. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 134–146, New York, NY, USA, 2003. ACM Press.

[10] R. Dube, C. Rais, K. Wang, and S. Tripathi. Signal stability based adaptive routing (ssa) for ad hoc mobile networks, 1997.

[11] Jabed Faruque and Ahmed Helmy. RUGGED: RoUting on finGerprint Gradients in sEnsor Networks. In *Proceedings of the IEEE International Conference on Pervasive Services (ICPS)*, Novi Sad, Serbia and Montenegro, July 2004.

[12] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly Resilient, Energy-efficient Multipath Routing in Wireless Sensor Networks. *ACM Mobile Computing and Communications Review*, 5(4), October 2001.

[13] Zygmunt J. Haas and Marc R. Pearlman. The performance of query control schemes for the zone routing protocol. In *SIGCOMM*, pages 167–177, 1998.

[14] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. In *Proceedings of the ACM/IEEE Internation Conference on Mobile Computing and Networking (Mobicom '00)*, Boston, USA, 2000.

[15] J.-Y. Le Boudec and M. Vojnovic. Perfect Simulation and Stationarity of a Class of Mobility Models. In *IEEE Infocom*, 2005.

[16] S. Jain and S. Das. Exploiting Path Diversity in the Link Layer in Wireless Ad Hoc Networks. In *Proceedings of the IEEE WoWMoM Symposium*, Taormina, Italy, June 2005.

[17] D. Johnson and D. Maltz. Dynamic Source Routing in Ad Hoc Wireless Networks. In T. Imelinsky and H. Korth, editors, *Mobile Computing*, volume 353, pages 153–181. Kluwer Academic Publishers, 1996.

[18] David B Johnson and David A Maltz. Dynamic source routing in ad hoc wireless networks. In Imielinski and Korth, editors, *Mobile Computing*, volume 353. Kluwer Academic Publishers, 1996.

[19] P. Larsson. Selection Diversity Forwarding in a Multihop Packet Radio Network with Fading Channel and Capture. *SIGMOBILE Mobile Computer Communication Review*, 5(4):47–54, 2001.

[20] Vincent Lenders, Martin May, and Bernhard Plattner. Density-based vs. Proximity-based Anycast Routing for Mobile Networks. In *IEEE INFOCOM*, Barcelona, Spain, April 2006.

[21] Q. Li, M. D. Rosa, and D. Rus. Distributed Algorithms for Guiding Navigation across a Sensor Network. In *Proceedings of ACM MobiCom*, San Diego, California, September 2003.

[22] A. R. Mitchell and D. F. Griffiths. *The Finite Difference Method in Partial Differential Equations*. Wiley, New York, 1980.

[23] W. Navid and T. Camp. Stationary Distributions for the Random Waypoint Model. *IEEE Transactions on Mobile Computing*, 3(1):99–108, 2004.

[24] E.F. Nogotov. *Applications of Numerical Heat Transfer*. Unesco, 757000 Paris, France, 1978.

[25] V. Park and S. Corson. *Temporally-Ordered Routing Algorithm (TORA)*. IETF Internet Draft, July 2001.

[26] C. Perkins. Ad-hoc on-demand distance vector routing. In *MILCOM '97 panel on Ad Hoc Networks*, 1997.

[27] Charles Perkins and Pravin Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, pages 234–244, 1994.

[28] C. K. Toh. Associativity based routing for ad hoc mobile networks. *Wireless Personal Communications Journal, Special Issue on Mobile Networking and Computing Systems*, 4(2):103 – 139, March 1997.

[29] S. Toumpis and L. Tassiulas. Packetostatics: Deployement of Massively Dense Sensor Networks as an Electrostatic Problem. In *IEEE INFOCOM*, Miami, USA, March 2005.

[30] A. Tsirigos and Z. Haas. Analysis of Multipath Routing - Part I: The Effect on the Packet Delivery Ratio. *IEEE Transactions on Wireless Communications*, 3(1), January 2004.

[31] Vincent Lenders and Martin May and Bernhard Plattner. Service Discovery in Mobile Ad Hoc Networks: A Field Theoretic Approach. In *Proceedings of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Taormina, Italy, June 2005.

[32] Xiang Zeng, Rajive Bagrodia, and Mario Gerla. Glomosim: A library for parallel simulation of large-scale wireless networks. In *Workshop on Parallel and Distributed Simulation*, pages 154–161, 1998.