

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

Vincent Lenders
Princeton University
vlenders@princeton.edu

Rainer Baumann
ETH Zurich
baumann@tik.ee.ethz.ch

Abstract—We present *link-diversity routing*, a routing paradigm that achieves high 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 towards the destination at the intermediate hops. This choice maximizes the opportunities to make progress at every hop in the presence of unpredictable link failures caused by mobility or fading effects. 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. We perform simulations of our algorithm with a DSDV-based implementation. Our simulations show that link-diversity routing increases the end-to-end packet delivery ratio to a factor of up to four without any additional protocol overhead compared to the traditional minimum hop-count based DSDV.

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 destination¹. 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 the routing protocol has converged to its new state), 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 robust as it works even in the presence of intermediate 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




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.

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.

Existing routing protocols are agnostic to the link diversity at the intermediate nodes and typically optimize for metrics like the number of hops [1], [2], [3], [4], the link qualities [5], [6], [7], or the expected throughput [8]. These approaches are effective in rather static networks where the link failure probability is low. However, as illustrated before, in wireless and mobile networks, ignoring the link-diversity often leads to packet drops caused by the lack of forwarding opportunities at the intermediate nodes.

In this paper, we propose *link-diversity routing*, a novel routing paradigm that increases the path resilience in the presence of unpredictable 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 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.

¹A 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.

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.

A. Heat-based Routing Paradigm

The key idea of link-diversity routing is to choose 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. From physics, we know that 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:

$$\Delta\phi = 0, \quad (1)$$

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 implies 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. In principle, the temperature of the source and destination could be chosen differently as long as $\phi(x_d) > \phi(x_s)$. These two boundary conditions together with the heat

²where 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.

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, i.e., a node which is not a heat source and has temperature that is higher than any of its neighbors. This is a necessary condition for the routing to work as we will see later.

Figure 3 shows the temperature field that follows the heat equation 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). The temperature of the intermediate nodes were obtained using the finite difference method as we will describe later.

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 (this follows directly from the definition of the field as given in Equation (1)), 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, representing an ascending gradient in that direction. In our approach however, unless a link along the steepest gradient is broken, we always forward packets along the steepest ascending gradient (represented with a thick arrow in the figure), corresponding to the neighbor with the 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 steepest gradient path towards the bottom is more resilient to link failures than the upper path.

Another interesting property of the routing model is 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.

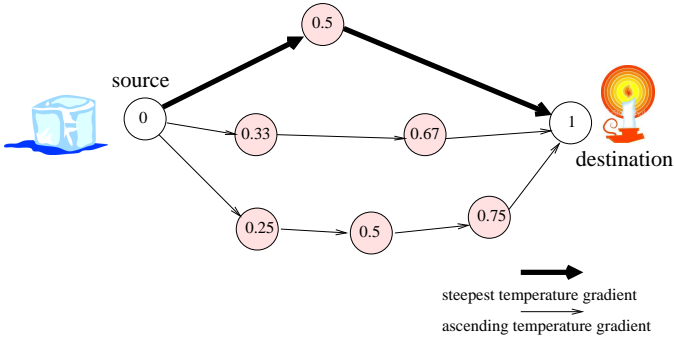


Fig. 4. Example temperature field: All paths have the same link diversity. The steepest gradient is therefore along the shortest path.




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

B. The FDMR Routing Algorithm

Establishing routes with our heat-based 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: (i) The algorithm should be completely distributed, and (ii) every node should be able to calculate its own temperature locally, based on the temperature of *only* its direct neighbors. Designing an algorithm that fulfills both requirements simultaneously is not trivial. We propose the FDMR algorithm. The FDMR algorithm relies on a numerical technique called the *finite difference method* [9] 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 temperature at a point is equal to the arithmetic average of the temperatures on a boundary surrounding this point. For example, in a 2-dimensional square grid, the temperature at a point is equal to the average temperatures of the four immediately neighboring points in the grid (see Figure 5). The way the finite difference method works is by iteratively approximating the temperature of the points on the grid as the arithmetic average of the neighboring points until the temperatures 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




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

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.

We define the FDMR algorithm as follows. Let $X = \{x_1, \dots, x_n\}$ 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 and have a fixed temperature) start with an initial temperature value $\phi_{t=0}(x_i) = 0; x_i \in X \setminus \{x_s, x_d\}$ 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 except the destination have a temperature value of 0. At each next iteration step, the nodes recalculate their temperature according to the new temperature.

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 routing protocol. Note however that the FDMR routing algorithm could as well be implemented using a link state routing protocol, but we do not further investigate this at this point. For our implementation, we used the DSDV [3] routing protocol, a pro-active protocol designed for wireless ad hoc networks. In the following, we

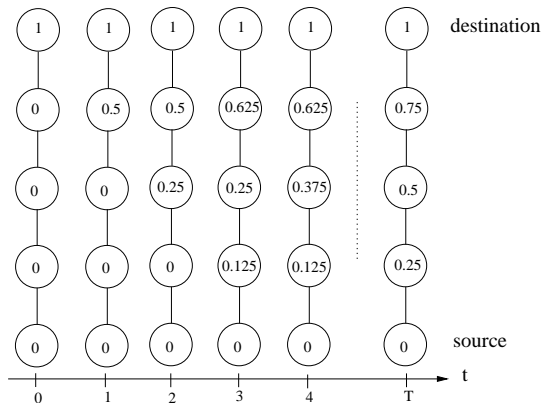


Fig. 7. Basic operation of the FDMR algorithm. At iteration step $t = T$, the algorithm has converged and the temperature distribution follows Equation (1).

describe how we adapted the DSDV routing protocol and how packet forwarding is done according to temperatures.

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 its 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 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 packets along the steepest ascending gradient corresponding to the neighbor with the highest temperature.

When a link breaks on a used path, 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 are neighbors with larger temperatures, 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 converged to the new topology, or dropped if the cache is full before an update has occurred.

IV. EVALUATION

To evaluate the benefits of link-diversity routing, we compare our DSDV-based implementation of the FDMR routing algorithm with the traditional minimum hop-count DSDV protocol in two specific ad hoc networking scenarios. The purpose of this comparison is to show how link-diversity routing improves the robustness compared to a routing protocol that uses the minimum hop-count to reach the destination. Since DSDV is a pro-active routing protocol, the overhead of the control messages is identical for both the minimum hop-count and our FDMR version, and is hence no further considered.

1) *Simulation Setup:* We use the Glomosim simulator [10] for our study. We consider the performance in a city-wide multi-hop ad hoc network in which the nodes are mobile and communicating over WLAN. The whole communication is infrastructure-less, i.e., there are no fixed access points and the mobile nodes act as relay/routers. These types of networks could be used for many different purposes in the area of urban inter-vehicle communication or person-to-person communication where a communication infrastructure is not available, damaged, or simply too expensive to use.

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 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 [11] on a network of streets and the random waypoint mobility model [2]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 which we extracted from a geographic information system (GIS) for the city of Zurich, Switzerland. In the random waypoint mobility model, the nodes move from randomly chosen waypoints in the square to other waypoints 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 mobility 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

³The random waypoint mobility model has shown to have non-desired behavior [12] when not well parameterized. We follow the guidelines as proposed in [13] to avoid such effects.

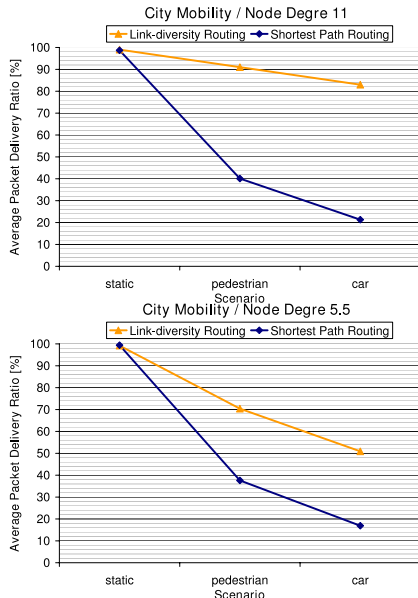


Fig. 8. Street mobility model: DSDV with traditional minimum hop-count compared to FDMR.

[10 – 20]m/s. These speeds corresponds to typical pedestrian and car speeds in a 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 Zurich are shown in Figure 8. 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 of the DSDV protocol 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 density increases. This is different for the FDMR algorithm. With the FDMR algorithm, the performance gets better as the average node degree increases. This is because the FDMR algorithm is able to exploit the link-diversity which becomes larger when the node degree increases.

The simulation results from the random waypoint model are given in Figure 9. The trends we observe with this mobility model are the same as with the previous mobility model. Overall, the packet delivery ratio is slightly better with this

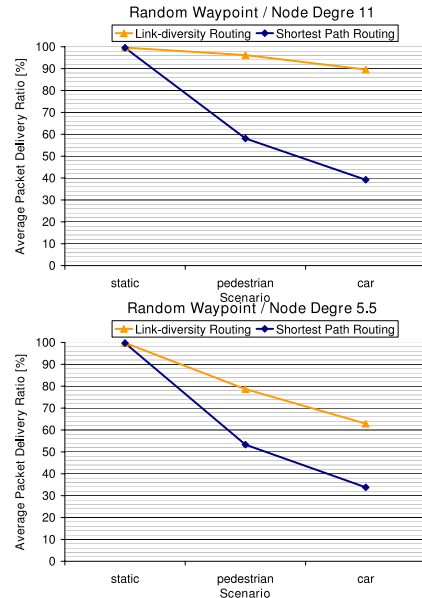


Fig. 9. Random waypoint mobility model: DSDV with traditional minimum hop-count compared to FDMR.

mobility model. The main reason is that the node distribution is not uniform. Nodes tend to concentrate around the center of the square [12]. 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 for metrics like the number of hops [1], [2], [3], [4], the link quality [5], [6], [7], or the expected throughput [8].

The authors of [14], [15], [16], [17] exploit the link-diversity in multi-hop routing to cope with lossy links and fading effects. These approaches 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. In contrast, we exploit the link-diversity to forward packets to nodes which have a higher number of outgoing links in order to increase the probability of successful packet delivery in the presence of unpredictable link failures.

Braided multipath routing [18] identifies multiple routes, using one as a primary and switching if the primary fails. Opportunistic multipath scheduling [19] splits traffic over multiple paths, adaptively favoring paths that provide low delays. Tsirigoris and Haas [20] propose to use erasure coded fragments of each packet over disjoint paths in a mobile ad hoc network, in order to tolerate losses 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 [21], [22], load balancing in the Internet [23], data collection in sensor networks [24], [25], sensor node placement [26], guided navigation [27], or service discovery [28]. 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 to achieve link-diversity routing, a property not addressed by those previous works. We first came up with the idea to use heat for anycast routing in [29]. This work differs by applying heat for unicast routing and by more rigorously applying the original model from physics.

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 unpredictable link failures caused by mobility or fading than traditional routing schemes. We show that by modeling 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 an adapted DSDV routing protocol, 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 additional routing protocol messages by changing only how the nodes compute their distance to the destination and hence with the same protocol overhead.

REFERENCES

- [1] Charles E. Perkins, Elizabeth M. Belding-Royer, and Samir R. Das, "Ad Hoc On-Demand Distance Vector (AODV) Routing," IETF Internet Draft, draft-ietf-manet-aodv-12.txt, November 2002.
- [2] D. Johnson and D. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," in *Mobile Computing*, T. Imelinsky and H. Korth, Eds., vol. 353, pp. 153–181. Kluwer Academic Publishers, 1996.
- [3] C. E. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers," in *Comp. Commun. Rev.*, October 1994, pp. pp 234–44.
- [4] Thomas Clausen and Philippe Jacquet, "Optimized Link State Routing Protocol," IETF Internet Draft, draft-ietf-manet-olsr-11.txt, July 2003.
- [5] C. K. Toh, "Associativity based routing for ad hoc mobile networks," *Wireless Personal Communications Journal, Special Issue on Mobile Networking and Computing Systems*, vol. 4, no. 2, pp. 103 – 139, March 1997.
- [6] R. Dube, C. Rais, K. Wang, and S. Tripathi, "Signal stability based adaptive routing (ssa) for ad hoc mobile networks," 1997.
- [7] 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)*, 2000, pp. 1697–1701.
- [8] Douglas S. J. De Couto, Daniel Aguayo, John Bricket, and Robert Morris, "A High-Troughput Path Metric for Multi-Hop Wireless Routing," in *Proceedings of ACM Mobicom*, San Diego, CA, USA, 2003.
- [9] A. R. Mitchell and D. F. Griffiths, *The Finite Difference Method in Partial Differential Equations*, John Wiley and Sons Ltd., New York, first edition, 1980.
- [10] Xiang Zeng, Rajive Bagrodia, and Mario Gerla, "GloMoSim: A Library for Parallel Simulation of Large-scale Wireless Networks," in *Proceedings of the 12th Workshop on Parallel and Distributed Simulations (PADS '98)*, Banff, Alberta, Canada, May 1998.
- [11] J.-Y. Le Boudec and M. Vojnovic, "Perfect Simulation and Stationarity of a Class of Mobility Models," in *IEEE Infocom*, 2005.
- [12] 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*, vol. 2, no. 3, pp. 257–269, 2003.
- [13] W. Navid and T. Camp, "Stationary Distributions for the Random Waypoint Model," *IEEE Transactions on Mobile Computing*, vol. 3, no. 1, pp. 99–108, 2004.
- [14] Sanjit Biswas and Robert Morris, "ExOR: Opportunistic Multihop Routing for Wireless Networks," in *Proceedings of ACM SIGCOMM'05*, Philadelphia, PA, USA, August 2005.
- [15] P. Larsson, "Selection Diversity Forwarding in a Multihop Packet Radio Network with Fading Channel and Capture," *SIGMOBILE Mobile Computer Communication Review*, vol. 5, no. 4, pp. 47–54, 2001.
- [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] R. Roy Choudhury and N. Vaidya, "MAC Layer Anycasting in Wireless Networks," in *Proceedings of the HotNets II Workshop*, Cambridge, MA, USA, November 2003.
- [18] 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*, vol. 5, no. 4, October 2001.
- [19] C. Cetinkaya and E. Knightly, "Opportunistic Traffic Scheduling over Multiple Network Paths," in *Proceedings of IEEE INFOCOM*, Hong Kong, September 2004.
- [20] A. Tsigros and Z. Haas, "Analysis of Multipath Routing - Part I: The Effect on the Packet Delivery Ratio," *IEEE Transactions on Wireless Communications*, vol. 3, no. 1, January 2004.
- [21] V. Park and S. Corson, *Temporally-Ordered Routing Algorithm (TORA)*, IETF Internet Draft, July 2001.
- [22] Vincent Lenders, Martin May, and Bernhard Plattner, "Density-based vs. Proximity-based Anycast Routing for Mobile Networks," in *IEEE INFOCOM*, Barcelona, Spain, April 2006.
- [23] Anindya Basu, Alvin Lin, and Sharad Ramanathan, "Routing Using Potentials: A Dynamic Traffic-Aware Routing Algorithm," in *Proceedings of ACM SIGCOMM*, Karlsruhe, Germany, August 2003.
- [24] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," in *Proceedings of ACM Mobicom*, Boston, USA, 2000.
- [25] Javed 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.
- [26] S. Toumpis and L. Tassiulas, "Packetostatics: Deployment of Massively Dense Sensor Networks as an Electrostatic Problem," in *IEEE INFOCOM*, Miami, USA, March 2005.
- [27] 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.
- [28] Vincent Lenders, Martin May, and Bernhard Plattner, "Service Discovery in Mobile Ad Hoc Networks: A Field Theoretic Approach," in *Proceedings of IEEE WoWMoM*, Taormina, Italy, June 2005.
- [29] Rainer Baumann, Simon Heimlicher, Vincent Lenders, and Martin May, "HEAT: Scalable Routing in Wireless Mesh Networks using Temperature Fields," in *Proceedings of IEEE WoWMoM*, Helsinki, Finland, June 2007.