

Density-based Anycast: A Robust Routing Strategy for Wireless Ad Hoc Networks

Vincent Lenders, Martin May, *Member, IEEE*, and Bernhard Plattner, *Member, IEEE*

Abstract—Existing anycast routing protocols solely route packets to the closest group member. In this paper, we introduce density-based anycast routing, a new anycast routing paradigm particularly suitable for wireless ad hoc networks. Instead of routing packets merely on proximity information to the closest member, density-based anycast routing considers the number of available anycast group members for its routing decision. We present a unified model based on potential fields that allows for instantiation of pure proximity-based, pure density-based, as well as hybrid routing strategies. We implement anycast using this model and simulate the performance of the different approaches for mobile as well as static ad hoc networks with frequent link failures. Our results show that the best performance lies in a tradeoff between proximity and density. In this combined routing strategy, the packet delivery ratio is considerably higher and the path length remains almost as low than with traditional shortest-path anycast routing.

Index Terms—Routing, Protocols, Mobile communication, Wireless communication, Anycast.

I. INTRODUCTION

Anycast (e.g., IP anycast [1]) is an addressing mode in which the same address is assigned to multiple hosts. Together, these hosts form an anycast *group* and each host is referred to as an anycast group *member*. Packets from a client destined to the group address are routed to the anycast group member closest to the client, where "closest" is in terms of the metrics used by the underlying routing protocol.

The most prominent use of anycast today is in the Internet to find replicated DNS root servers [2] or to locate rendezvous points in multicast trees [3]. However, anycast has also many potential applications in wireless ad hoc networks. For example, anycast can be used in wireless mesh networks to route data packets to an Internet gateway, or in sensor networks to send data to any data sink when multiple sinks are accessible.

Today's anycast routing protocols are most commonly modifications of existing unicast routing protocols. For example, link-state routing protocols such as OSPF [4] have been extended to support anycast routing by adding a virtual node that represents the anycast service [5]. With distance vector routing algorithms such as RIP [6], anycast routing is implemented by group members that advertise their anycast address with a distance of zero [5]. Also in the context of ad hoc networks, link reversal algorithms such as TORA [7] were extended to support anycast routing by assigning a height of zero to all members of a given anycast group [5]. Since these proposed anycast protocols are designed as extensions of unicast routing techniques, they are easy to implement and to deploy. However, as a consequence, they all follow the routing strategy determined by the corresponding unicast

routing technique: packet delivery to the closest group member using shortest-path forwarding.

In this paper, we propose a new method that adds a whole family of routing strategies to the class of anycast routing schemes: density-based routing. This method not only considers the member proximity in the routing decision, but also the quantity of accessible group members. Therefore, it is possible that a path over which multiple members are accessible is preferred over a path to a closer single member. Our goal is not to replace proximity-based routing, but to add a new dimension to the routing strategy design space. With this new axis in the design space, the routing strategy can be designed as an optimal tradeoff between proximity and density.

To assess merits of this novel idea, we develop a model that incorporates both strategies. Our model is based on potential fields from physics and allows using a single parameter to prioritize proximity or density in the routing decision. It is for example possible to model the behavior of traditional anycast routing algorithms that always select the shortest path to the closest group member. On the other hand, it is possible to base the routing decisions on a tradeoff between proximity and density, or solely on density where the distance to the group members is not considered at all in the routing decision.

Density-based anycast routing is of particular interest in wireless and mobile ad hoc networks where the network conditions are dynamic. Proposed anycast routing protocols for such networks [5], [8], [9], [10], [11] are all implemented as modifications of existing unicast routing protocols and hence always route packets towards the closest group member. In wireless and mobile networks however, the closest member might become unreachable because it moves away or because of intermediate link failures along the path. Under such conditions, routing towards a dense group member population increases the probability that a packet eventually reaches any group member because packets can more easily be re-routed to neighboring group members when the targeted one becomes unreachable.

To evaluate the benefits of density versus proximity anycast routing, we implement a distributed protocol to establish potential fields and simulate the different routing strategies under various dynamic network conditions. In particular, we consider two different mobile scenarios and a static scenario with temporary link failures. Our results show that a combined proximity-density routing strategy increases considerably the robustness (in terms of packet delivery ratio) compared to the traditional pure proximity-based routing scheme while not increasing significantly the path length of the traveled packets.

The rest of this paper is organized as follows. In Section II,

we present our anycast routing model and evaluate it in Section III. In Section IV, we introduce a simple distributed routing protocol for routing based on potential fields. This protocol is then used in Section V to compare pure proximity, pure density and combined routing strategies. We describe related work in Section VI and conclude in Section VII.

II. ANYCAST ROUTING MODEL

In this section, we present our field-based model for anycast routing. First, we give an overview of the basic concepts. Then, we introduce our definition of potential fields in a networking context and describe how packets are routed along those fields. Finally, we discuss convergence limitations of the model due to undesired local maxima in the fields that might occur in particular network topologies.

A. Overview

Our model is inspired from field theory in physics. Every group member creates a potential field which decreases with d^{-k} , where d is the distance to the group member, and k determines how quickly the field decreases. The field of an entire anycast group is the linear superposition of all individual fields from the group members. An example field for an anycast group with four members (marked as black nodes) is depicted in Figure 1. The peaks in the field are at the locations of the group members. Note that only one field is drawn in the figure, but as each anycast group requires its own field, multiple fields will generally co-exist simultaneously.

We achieve anycast routing by forwarding packets towards the steepest gradient of the field. This is in analogy to field diffusion in physics. By following the steepest gradient, packets eventually reach a field maximum, i.e., a group member. The steepest gradient at each node is determined by comparing the potential values φ of its neighbors. The steepest gradient is towards the neighbor with the highest potential value.

The proposed routing model allows for different anycast strategies comprising proximity, density, and combined routing strategies. Which routing strategy is applied is determined by the value of the parameter k . We will show in the next section that a proximity-based routing strategy (the routing strategy of existing anycast routing protocols which consists of forwarding packets to the closest member along the shortest path) is modeled by setting $k > \mu$, where μ is a constant depending on the network size and the anycast group size. We also show that for $0 < k \leq \epsilon$ (where $\epsilon < \mu$), a pure density-based routing strategy is modelled where the proximity of the group members is no longer considered for routing decisions. By choosing a value for k between μ and ϵ , we are able to model combinations of these two one-sided routing strategies.

B. Potential fields

We define the potential field of an anycast group *member* j with a strictly decreasing function. That is, we define the potential value at some node n as

$$\varphi_j(n) = \frac{1}{d_j^k(n)}, \quad 0 < k < \infty \quad (1)$$




Fig. 1. Example potential field. Black nodes represent group members.

where $d_j(n)$ is defined as the distance of n to the group member j , and the exponent k is a parameter that determines how quickly the potential decreases with increasing distance to the group member. In this paper, we always use the number of hops to calculate the distance between a node and a group member. However, the distance function could also capture different metrics (such as for example the transmission delay between a node and a group member) as long as the potential function remains strictly decreasing.

Anycast groups consists of multiple members. Every member contributes to the field of the group. Thus, the potential field of an anycast *group* \mathcal{N} is defined as the superposition of the potential fields of all members in this group:

$$\varphi(n) = \sum_{j \in \mathcal{N}} \varphi_j(n) = \sum_{j \in \mathcal{N}} \frac{1}{d_j^k(n)} \quad (2)$$

With this definition, the potential field's shape resembles a landscape with poles at every group member $\forall j \in \mathcal{N}$ since $\varphi \rightarrow \infty$ (in one term of Equation (2) the distance $d_j(j)$ is equal to zero hops). By varying the exponent k , the steepness of the field varies: the larger the value of k , the steeper the field.

C. Gradient-based routing

We use the potential field of the anycast group to route packets in the network. The routing mechanism is similar to field diffusion in physics. With field diffusion, an element (e.g., a test charge in an electrical field) is attracted by a force in the direction of the steepest gradient of the field. If the element is free to move, it will diffuse along the steepest gradient until it arrives at a field maximum. In the same manner, we route anycast packets along the steepest gradient of the potential field. The steepest gradient at each node is determined by evaluating the potential values of the neighbors. That is, the link from a node to the neighbor with the highest potential value corresponds to the steepest gradient. Therefore, the nodes always compare the potential value of their neighbors and forward anycast packets to the neighbor with the highest potential value. A necessary condition is that the potential at the neighbor with the highest potential value is larger than at

the forwarding node. This guarantees that the steepest gradient is ascending.

D. Convergence Limitations

The proposed model manages to successfully deliver anycast packets from any node in the network with the condition that there are no local¹ maxima in the potential field. In a local maximum, packets are stuck since there are no neighbors with a greater potential value to forward to. However, in contrast to the physical model which operates in a continuous space, we cannot avoid the occurrence of local maxima in our field model which operates on a discrete set of nodes. This fundamental difference implies that we can only guarantee the absence of local maxima for $\infty > k > c$, where c is a constant we will derive in the next section. For values of $k \leq c$, local maxima may occur and routing along the steepest gradient may not converge. Although we cannot guarantee routing convergence for that range, we use our model in that range because (i) the occurrence of local maxima in random networks is rare (see Section III) and (ii) in practice, local maxima can easily be detected by comparing the own potential value with the value of the direct neighbors. Hence, it is possible to circumvent them at the protocol layer without considerable communication overhead.

III. MODEL EVALUATION

The goal of this section is to evaluate the routing model and to provide a more profound understanding on how the variable k affects the routing strategy. The model evaluation comprises two parts: (i) a proof that field-based routing is loop-free and (ii) an analysis of the convergence properties of the model. The analysis in this section assumes static network conditions; the effect of node mobility and network dynamics is analyzed in Section V.

A. Loop-freeness

An important characteristic of a routing algorithm is to be loop-free. We present a sketch of a proof for the following theorem for loop-free routing.

Theorem 1: Anycast routing along the steepest ascent in a potential field as defined in Equation (2) is loop-free for any $k > 0$.

Sketch of Proof: Gradient-based routing along the steepest ascent requires that the potential value at every hop is strictly larger than at the previous hop. In a loop, a packet would have to traverse one node more than once which is not possible per definition. ■

B. Convergence of gradient-based routing

1) *Convergence bound for k :* By following the steepest gradient, packets from any node in a connected network will always reach a group member unless there is a local maximum in the field. We next derive an upper bound for k , which

¹In this paper, we refer to local maxima as undesired maxima in the field at nodes that are not group members. Global maxima are desired field maxima at the group members.

depends on both the network and group size, where we can guarantee that no local maxima exist in the potential field.

Lemma 2: Consider a potential field in a connected network with diameter² D created by an anycast group of size N . The potential field shows no other maxima than at the nodes which are group members if k is any constant satisfying:

$$k > \frac{\log N}{\log \frac{D}{D-1}} \quad (3)$$

Proof: Consider any node x which is not member of the group. To guarantee that the potential value at x is not a maximum, there must at least be one neighbor y with a greater potential value:

$$\varphi(x) < \varphi(y) \quad (4)$$

We denote \mathcal{N} as the complete set of group members and assume that the distances from x to all group members ($d_i, \forall i \in \mathcal{N}$) are known. Then, we can express $\varphi(x)$ and $\varphi(y)$ as

$$\begin{aligned} \varphi(x) = \sum_{i \in \mathcal{N}} \frac{1}{d_i^k} &< \sum_{i \in \mathcal{N}_1} \frac{1}{(d_i - 1)^k} + \sum_{i \in \mathcal{N}_2} \frac{1}{d_i^k} \\ &+ \sum_{i \in \mathcal{N}_3} \frac{1}{(d_i + 1)^k} = \varphi(y), \end{aligned} \quad (5)$$

where $\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3$ are disjoint subsets of \mathcal{N} ($\mathcal{N}_1 \cup \mathcal{N}_2 \cup \mathcal{N}_3 = \mathcal{N}$ and $\mathcal{N}_u \cap \mathcal{N}_v = \emptyset, \forall u, v = 1, 2, 3; u \neq v$). The potential value of y must be of this form since x and y are direct neighbors, and the distance of y to any group member can only be one hop smaller, equal, or one hop larger than the distance of x to the corresponding member. Since all terms in Equation (5) are positive by definition, if the potential value at y is still larger than the potential value at x when only considering the contribution from the group members \mathcal{N}_1 , the potential at node x cannot be a local maximum. Thus, we can simplify the previous condition as

$$\varphi(x) = \sum_{i \in \mathcal{N}} \frac{1}{d_i^k} < \sum_{i \in \mathcal{N}_1} \frac{1}{(d_i - 1)^k} \leq \varphi(y). \quad (6)$$

By only considering the contribution at y from the closest group member s , if the potential at x is still smaller than at y , there is no local maximum at x :

$$\varphi(x) = \sum_{i \in \mathcal{N}} \frac{1}{d_i^k} < \frac{1}{(d_s - 1)^k} \quad (7)$$

We now consider the worst-case potential value of x . The worst case value is when the potential is maximal. The maximal potential value for x is obtained when the distances to the group members are minimal. Since we said that s is the closest member, d_s is the smallest distance that any group member can have. Therefore, the worst case is when all group members are at distance d_s to node x :

$$\varphi(x) \leq \frac{N}{d_s^k} < \frac{1}{(d_s - 1)^k} \quad (8)$$

²Note that the network diameter D in this paper is defined as the longest shortest path between any two nodes in the network.




Fig. 2. Example of a local maximum for $k = 1$. Gray nodes are group members.

We now solve for k and get

$$k > \frac{\log N}{\log \frac{d_s}{d_s-1}} \quad (9)$$


The distance d_s is strictly smaller than the network diameter D and since $\log \frac{d_s}{d_s-1} > \log \frac{D}{D-1}$, the condition that the potential value of x is smaller than the potential of any neighbor node y is

$$k > \frac{\log N}{\log \frac{D}{D-1}}. \quad (10)$$

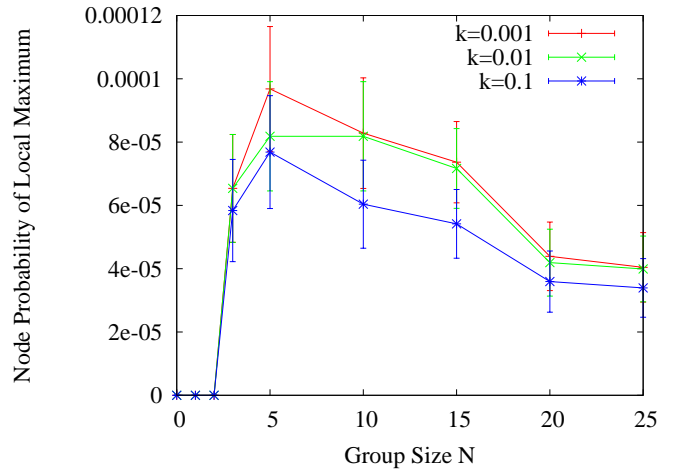
■

We identified the range of k where we can guarantee that there are no local maxima and thus guarantee that routing converges. However, as we will see later, density-based routing strategies require a smaller value of k . Consequently, we also use our model for smaller values of k than the derived bound. When k is smaller than the derived bound, there is no strict guarantee that a packet will reach its destination since local maxima may occur in the potential field. We next look at how frequent such local maxima occur in random network topologies.


2) *Local maxima in random topologies*: Local maxima may occur in particular topologies like for example the one shown in Figure 2. The node in the center of the star-shaped topology has a potential value of 2.5 which is higher than the potential value of any of its surrounding neighbors. To assess how frequent local maxima occur for $k \leq \frac{\log N}{\log \frac{D}{D-1}}$, we perform simulations with random network topologies. Our simulations are based on graphs which are constructed as follows. We place a number of nodes randomly on a square. We assign a link between two nodes in the graph if the geometric distance between those two nodes is smaller than the wireless range, a constant value equal for all nodes. The distances of the nodes to the group members, as required to compute the potential values, are obtained by computing the shortest path according to Dijkstra's algorithm [12] in the graph. Finally, to detect whether there is a local maximum in the potential field, we iterate over all nodes and compare their potential values with the values of their direct neighbors in the graph. If a node's



(a) Probability vs. k .



(b) Probability vs. group size N .



(c) Probability vs. network density

Fig. 3. Local maximum node probability in random networks.

potential is larger than all of its neighbors, and it is not itself a group member, the node constitutes a local maximum.

Figure 3(a) shows the probability that a non-member node becomes a local maximum for values of k between 1000 and 10^{-5} in graphs of 500 nodes. The used networks have an average node degree of approximately 12. The error-bars in the plot represent a 90 percent confidence interval. For large values of k ($k > 3$), there are no local maxima in the potential field independent of the group size. For $k < 3$, some very few local maxima form. However, the probability remains quite low; the probability is always below 0.01% for the used group sizes.

We also plot the probability of local maxima versus the group size on the horizontal axis using the same parameters in Figure 3(b). For group sizes of 2 and below, there are no local maxima. Local maxima occur for group sizes larger than 3. Note, that the highest probability for a local maximum is reached with a group size of 5. For larger groups, the probability decreases significantly. This is mainly because local maxima form only in specific star topologies, when the group members are in different directions but about the same distance from a node. By increasing the group size, the probability that such a topology occurs, decreases.

In Figure 3(c), the occurrence of local maxima is plotted for different network densities. We express the network density as the average node degree of the nodes on the horizontal axis. We obtained networks of different average node degrees (densities) by varying the number of nodes while keeping the simulation area constant. For the smallest average node degree of 7.8, the number of nodes was 250 and for the largest average node degree of 20.5, the number of nodes was 700. We observe a higher probability for local maxima in sparse networks. For example, with $k = 0.1$, the probability of a local maximum at a node is 0.018% for a node degree of 7.8. However, when increasing the node degree to 20.5, the probability decreases to 0.0027%. This result confirms our expectation. By increasing the node density in our model, the discrete field approaches a continuous field which does not show any local maxima. As previously mentioned, for $k \geq 3$, there are no local maxima in the field.

C. Effect of k on routing strategy

The potential field and thus also the resulting routing strategy is influenced in our model by the exponent k in Equation (1). For large values of k , the field results in a steeper distribution than for small values. In the following, we derive a bound for k where packets are always routed towards the closest group member along the shortest path (proximity routing), and a bound for k where packets are only routed in a specific direction based on the group member density and the member proximity is irrelevant (pure density routing).

1) Proximity-based routing:

Lemma 3: Consider a potential field in a connected network with diameter D created by an anycast group of size N . A packet from any node in the network is always routed to the closest group member of the anycast group along the shortest




Fig. 4. Network topology with one member at distance d_s and $N - 1$ members at distance $d_l > d_s$.

path, if k is a constant satisfying

$$k > \mu(N, D) \quad (11)$$

(see Table I for sample values of μ)

Proof: We prove this relation by determining the steepest gradient in a worst-case scenario and show that it points to the shortest path. The worst case scenario is obtained when all (except one) members of a group are located in the opposite direction of the closest group member. This scenario is depicted in Figure 4. The node m_s is the closest group member to the sender, however, the remaining $N - 1$ group members (m_{l1}, \dots, m_{lN}) are all located at the opposite end. This arrangement is the worst case because m_{l1}, \dots, m_{lN} create the largest possible potential at node n_2 which is in the opposite direction of the closest member, and the smallest possible potential at n_1 which is on the shortest path to m_s . If we can guarantee that $\varphi(n_1)$, which is the potential of the node on the shortest path to m_s , is larger than $\varphi(n_2)$, then a packet following the steepest ascent will be forwarded to n_1 towards m_s over the shortest path. Therefore, the following condition must be met:

$$\varphi(n_1) > \varphi(n_2)$$

Without loss of generality, we assume that m_{l1}, \dots, m_{lN} are all equidistant with d_l from the sender and we calculate the potential values using Equation (2):

$$\varphi(n_1) = \frac{1}{(d_s - 1)^k} + \sum_{N-1} \frac{1}{(d_l + 1)^k} > \frac{1}{(d_s + 1)^k} + \sum_{N-1} \frac{1}{(d_l - 1)^k} = \varphi(n_2) \quad (12)$$

The worst-case for this inequality is obtained when d_l is small and d_s is large. Therefore, we set the distance d_l to the smallest possible distance $d_l = d_s + 1$ (which is one hop more than the distance to the closest member m_s) and d_s to the largest possible value $d_s = \frac{D-1}{2}$ (since $D = d_s + d_l$) and get

$$\frac{1}{\left(\frac{D-1}{2} - 1\right)^k} + \frac{N-1}{\left(\frac{D-1}{2} + 2\right)^k} > \frac{1}{\left(\frac{D-1}{2} + 1\right)^k} + \frac{N-1}{\left(\frac{D-1}{2}\right)^k}. \quad (13)$$

This equation cannot be solved analytically. We therefore bring it into a cancellation-free form to be able to solve it numerically (see Appendix in [13] for further details how we obtained this form):

$$f(k) > N - 1 \quad (14)$$

with

$$f(k) = \left(\frac{\frac{D-1}{2}(\frac{D-1}{2} + 2)}{(\frac{D-1}{2} + 1)(\frac{D-1}{2} - 1)} \right)^{\frac{k}{2}} \cdot \frac{\sinh\left(\frac{k}{2} \log\left(\frac{\frac{D-1}{2}-1}{\frac{D-1}{2}+1}\right)\right)}{\sinh\left(\frac{k}{2} \log\left(\frac{\frac{D-1}{2}}{\frac{D-1}{2}+2}\right)\right)}. \quad (15)$$

The solution μ for k with sample values of D and N is given in Table I. Since for $k > \mu$, the steepest gradient is always pointing to the shortest path, a packet from any node in the network is always routed along the shortest path. ■

2) *Pure density-based routing*: We now determine the other extreme in which packets are always forwarded towards the highest member density and the distance to the group members is irrelevant.

Lemma 4: Consider a potential field in a connected network with diameter D that is created by an anycast group of size N . A packet from any node in the network is routed independently of the distances to the members if k is any constant satisfying

$$0 < k \leq \epsilon(D, N) \quad (16)$$

and

$$N > \frac{\log 3}{\log \frac{D-1}{D-3}} + 1 \quad (17)$$

(see Table II for sample values of ϵ)

Proof: We prove this by calculating the steepest gradient in a best-case scenario and show that it does not point towards the next hop over the shortest path of the closest member. We assume again a network topology as depicted in Figure 4 where the closest group member m_s is at distance d_s from a sender. We determine the smallest exponent k when a packet is not routed towards the closest group member independent of how close it is to this member. If a packet is not routed along the shortest path, then the potential value of the next hop towards this member must be smaller or equal to at least one other neighbor:

$$\varphi(n_1) \leq \varphi(n_2) \quad (18)$$

or

$$\frac{1}{(d_s - 1)^k} + \sum_{N-1} \frac{1}{(d_l + 1)^k} \leq \frac{1}{(d_s + 1)^k} + \sum_{N-1} \frac{1}{(d_l - 1)^k}. \quad (19)$$

In the best case, m_s is very close to the sender and all other group members are far away. Therefore, we set d_s to the smallest possible value $d_s = 2$ and d_l to the largest possible value $d_l = D - d_s = D - 2$ (Note that we do not consider the trivial case $d_s = 1$ since then a packet will always be routed to m_s because it is now a direct neighbor of m_s with a potential of $\varphi \rightarrow \infty$). Now we get:

$$1 + \frac{N-1}{(D-1)^k} \leq \frac{1}{3^k} + \frac{N-1}{(D-3)^k}. \quad (20)$$

This inequality cannot be solved analytically and we therefore bring it in a cancellation-free form to solve it numerically (see Appendix in [13] for further details):

$$f(k) \leq N - 1 \quad (21)$$

with

$$f(k) = \left(\frac{(D-1)(D-3)}{3} \right)^{k/2} \cdot \frac{\sinh\left(\frac{k}{2} \log 3\right)}{\sinh\left(\frac{k}{2} \log\left(\frac{D-1}{D-3}\right)\right)}. \quad (22)$$

Note that this inequality has solutions only for $N - 1 > \frac{\log 3}{\log \frac{D-1}{D-3}}$. The solution ϵ of this inequality is given for sample values of D and N in Table II. Since for $k \leq \epsilon$, the steepest gradient does not correlate with the shortest path, packets are routed independently of the distance to the members. ■

3) *Combined density and proximity routing*: In contrast to pure proximity or pure density based anycast routing, our model allows for a whole family of combined routing strategies. These strategies follow from potential fields with a value of k between the previously identified bounds. To illustrate how the potential field (and also the routing) varies within this range, we investigate the shape of the fields on random graphs.

The random graphs we use are generated by placing nodes randomly on a square of 2000 by 2000 meters side length. There exists an edge between two nodes in the graph if the geometric distance between two nodes is less than a threshold value (corresponding to the wireless range of the devices). On these graphs, we determine the potential field created by randomly placed group members. As metric, we plot the ratio of nodes that changed their steepest gradient compared to the steepest gradient obtained with proximity-based routing. Precisely, the definition of our metric is the ratio of nodes in the network at which the neighbor with the largest potential value using a specific value for k is different from the neighbor with the largest potential value enforcing shortest-path routing. For example, a ratio of 0.5 for $k = 1$ indicates that at 50% of the nodes, the neighbor with the highest potential is different from the neighbor on the shortest path. Note that since this metric indicates how often the steepest gradient changes, it also captures the effect of changes in the routing strategy because packets are routed along the steepest gradient.

In Figures 5 and 6, we plot this metric for different group member sizes and average node degrees using a 90-percent confidence interval. The used graphs consist of 400 nodes. The resulting distribution allows for the distinction of the three different regimes. Pure proximity-based routing (marked with I) is the range of k when all packets follow the shortest path (ratio is zero). As soon as the ratio is non-zero (starting at $k \approx 100$), the routing decision is done as a combination of proximity and density routing. This regime (marked with II) ends when the ratio saturates at $k \approx 0.01$. For smaller values of $k \lesssim 0.01$, routing is purely density-based (marked with III) as further decreasing the value of k no longer changes the shape of the potential field.

In the combined proximity-density routing regime (II), we observe a non-static increase of the change ratio. By closely analyzing the results, we found that for $100 \gtrsim k \gtrsim 10$, the

D	6	8	10	12	14	16	18	20
4	1.795635	2.931176	4.049784	5.161040	6.268394	7.373423	8.476947	9.579444
6	2.936760	4.591675	6.224987	7.849220	9.468746	11.085512	12.700515	14.314324
8	3.657162	5.652006	7.623068	9.584353	11.540646	13.494042	15.445605	17.395938
10	4.183686	6.431778	8.654766	10.867593	13.075283	15.280017	17.482893	19.684523
12	4.598563	7.048652	9.472706	11.886338	14.294741	16.700152	19.103687	21.505970
14	4.940854	7.559027	10.150448	12.731245	15.306746	17.879221	20.449810	23.019141
16	5.232170	7.994312	10.729110	13.453120	16.171770	18.887369	21.601069	24.313505
18	5.485726	8.373796	11.234018	14.083305	16.927177	19.767971	22.606854	25.444466
20	5.710183	8.710176	11.681874	14.642505	17.597668	20.549730	23.499864	26.448721
22	5.911535	9.012257	12.084283	15.145120	18.200436	21.252628	24.302877	27.351841
24	6.094092	9.286393	12.449630	15.601562	18.747926	21.891138	25.032397	28.172357

TABLE I

SAMPLE VALUES OF μ . N IS THE NUMBER OF ANYCAST GROUP MEMBERS AND D IS THE NETWORK DIAMETER.

D	6	8	10	12	14	16	18	20
6	-	-	-	-	-	-	-	-
8	0.129472	-	-	-	-	-	-	-
10	0.343468	0.061210	-	-	-	-	-	-
12	0.511681	0.197150	0.031419	-	-	-	-	-
14	0.650004	0.309386	0.129192	0.015418	-	-	-	-
16	0.767310	0.404871	0.212506	0.090838	0.005733	-	-	-
18	0.869051	0.487902	0.285049	0.156559	0.066601	-	-	-
20	0.958808	0.561315	0.349261	0.214772	0.120540	0.050101	-	-
22	1.039066	0.627080	0.406841	0.267005	0.168956	0.095628	0.038259	-
24	1.111609	0.686620	0.459017	0.314358	0.212865	0.136930	0.077500	0.029415

TABLE II

SAMPLE VALUES OF ϵ . N IS THE NUMBER OF ANYCAST GROUP MEMBERS AND D IS THE NETWORK DIAMETER.

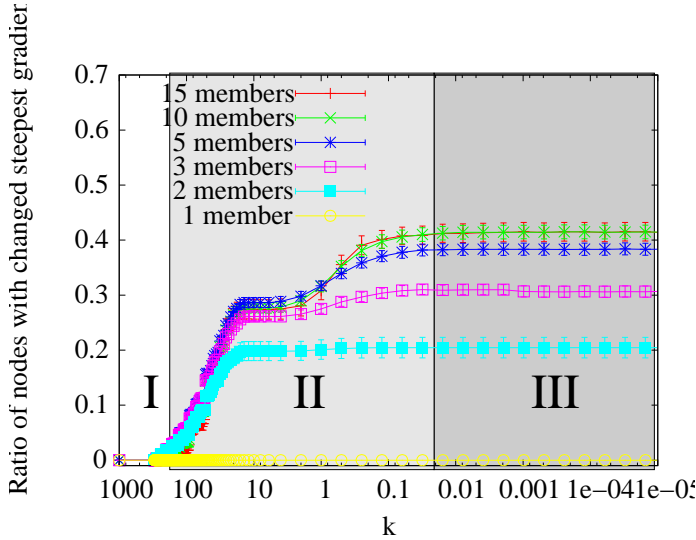


Fig. 5. Effect of k on the steepest gradient for varying group sizes.

steepest gradient changed, but the changes did not impact the path length from the nodes to the group members. Therewith, we conclude that there were more than one equidistant closest group members to a node. Thus, although density came into play, routing was still along the shortest path to the closest group member. For $0.01 \lesssim k \lesssim 10$, the changes of the steepest gradient had an impact on the path length and packets are no longer routed along the shortest path.

So far, we analyzed our model in a general sense on static graphs. To investigate how density-based routing performs in




Fig. 6. Effect of k on the steepest gradient for varying average node degrees.

wireless ad hoc networks, we propose in the next section a simple routing protocol to establish potential fields and forward packets along the steepest gradient. This protocol is then used in Section V in a simulation-based performance analysis including node mobility and link failures.

IV. FIELD-BASED ANYCAST ROUTING PROTOCOL

To evaluate the performance of density-based anycast routing in dynamic networks, we designed a distributed routing protocol to establish potential fields and forward packets along

the steepest gradient. Note that the focus of this paper is not on the performance of the routing protocol itself, but on the comparison of the different anycast routing strategies. Therefore, we designed a relatively simple protocol for the only purpose of comparing the different strategies, and leave possible enhancements of the protocol to future work.

A. Potential field establishment

To establish a potential field, every node in the network must know its distance to the existing group members. For this purpose, the group members periodically flood the network with a message indicating the anycast group they belong to, and their identity (i.e., an identifier that uniquely identifies the group member). These flooded messages also include a time-to-live (TTL) value indicating how many hops the packet has traveled. The TTL value serves two purposes. First, it allows every receiver to determine its distance to the group member who initiated the flooding. Second, it allows to limit the flooding scope by only rebroadcasting messages which have a TTL value greater than 0 which reduces the communication overhead produced by such messages. By listening to those messages, each node calculates its potential value according to Equation (2). The interval at which the member should advertise such messages is a tradeoff between accuracy and protocol overhead. In this paper, we do not focus on finding the best compromise with regard to this tradeoff, but we study the relative performance resulting from different anycast strategies using the same advertisement interval. Note that the impact of the TTL value on the performance of the routing is evaluated in Section V.

B. Gradient determination

To determine the steepest gradient of a field, the nodes in an ad hoc network must know the potential values of their direct neighbors. For this, neighboring nodes also exchange their potential values on a periodic basis. Such messages are one-hop broadcast packets and include a list of all the known anycast groups and the corresponding potential value for each group. Again, the rate at which such messages are exchanged is a tradeoff between accuracy and protocol overhead.

C. Packet forwarding

Packet forwarding is simply forwarding along the steepest gradient. Therefore, packets are forwarded to the neighbor with the highest potential value. If for any reason, the neighbor with the greatest potential value is unreachable (e.g., this neighbor might have moved away), the packet is simply forwarded to the neighbor with the second highest potential value. In case this node is also unreachable, the packet is forwarded to the neighbor with the next highest potential value, and so on. A node continues with this procedure until there are no neighbors left with a higher potential value than its own. Note that nodes are not allowed to forward to a neighbor with a lower potential value to make sure that routing eventually converges and loops do not form. If a node has no neighbors left with a larger potential value than its own, it drops the packet.

V. SIMULATION-BASED PERFORMANCE EVALUATION

To assess the performance of density-based anycast routing, we performed simulations with wireless ad hoc networks. Our aim is to study the robustness properties when the network conditions are dynamic. To this end, we first consider two mobile scenarios, one in which the nodes move according to the random waypoint mobility model [14]³, and another one in which the nodes move according to the random walk mobility model [17]. We also evaluate the performance for static wireless ad hoc networks where individual wireless links become temporarily unavailable according to a link error model proposed in [18]. Finally, we analyze the scalability of our approach by determining, how sensitive density-based routing is to the flooding scope of the member advertisement messages.

A. Simulation model

Our simulations are based on a simulator written in C we developed for the purpose of this work. Our simulator uses a simple communication model in which there exists a bi-directional link between two nodes if their geometric distance is smaller than the wireless range, a fixed value identical for all devices. We used an 802.11-like MAC layer where nodes can send a broadcast message to all neighbors or a unicast message to one specific neighbor. However, we do not model in details any physical layer and MAC-layer effects like interference, packet collisions, etc. Therefore, the only reason for an existing link to become unavailable is because two nodes move apart outside the wireless communication range (in the mobile scenarios) or when we explicitly declare a link unavailable (as we do in the static scenario). This simplified model allows us to study the cause of packet losses related to the routing protocol itself and not from link-layer transmission issues. We assume that the link-layer can recover from individual packet losses by means of retransmissions in the absence of a missing link-layer acknowledgement.

B. Robustness versus path stretch

In Figure 7, we plot the robustness and the path stretch averaged over at least 2000 simulation seeds of a network of 500 nodes which are moving according to the random waypoint mobility model [14]. The nodes all move with the same speed without pausing on a square of size 2500 by 2500 meters. The wireless range is set to 225 meters resulting in an average node degree of approximately 11 independent of the node speed. The routing strategies are compared for a group size of 10 members, which also move according to the random waypoint mobility model.

The robustness of each strategy is assessed versus the node speed in Figure 7(a), comparing the packet delivery ratio we obtain when random nodes are sending anycast packets. We used values of k ranging from $k = 0.01$ (pure density) to $k = 300$ (pure proximity). The results show that pure proximity-based routing performs worse as the node speed increases.

³The random waypoint mobility model has shown to have non-desired behavior when not well parameterized [15]. We follow the guidelines as proposed in [16] to avoid such effects.


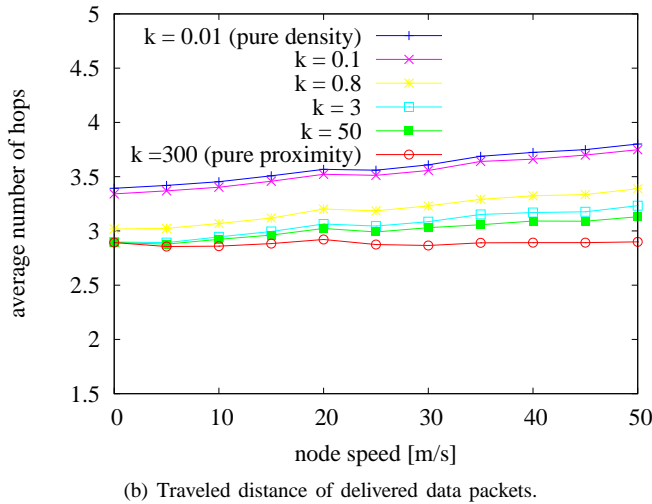




Fig. 7. Robustness and path stretch for random waypoint node mobility.

Fig. 8. Robustness and path stretch for random walk node mobility.

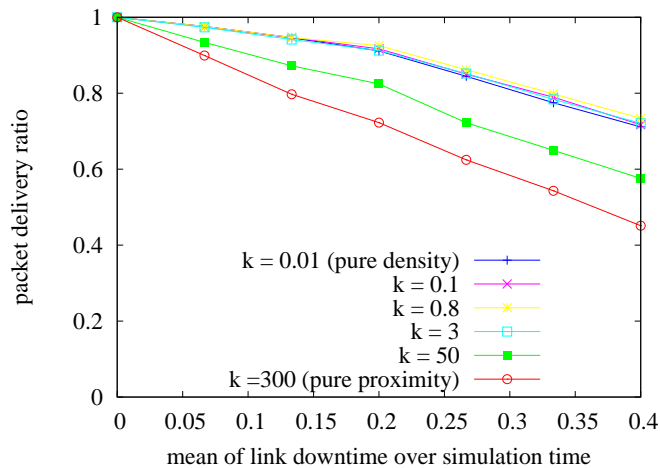
The smaller the value of k , the better the delivery ratio is. The superior performance of density-based routing strategies is based upon the increased probability to reach an alternative group member in the neighborhood when the targeted group member becomes unreachable due to node mobility.

Notice that from a value of $k = 3$ on, further decreasing k shows only a marginal performance improvement in the delivery ratio. To understand why, it is necessary to look at the traveled distance of packets as shown in Figure 7(b). The traveled distance is assessed as the average number of hops that anycast packets undertake until they reach a group member. As expected, lower values of k result in longer traveled distances as the routing strategies opt for member density instead of member proximity. However, when looking at the static case ($v = 0\text{m/s}$), we see that the path length increases only for $k \leq 0.8$. For $k = 3$ and $k = 50$, the traveled distance is identical to pure proximity routing ($k = 300$). Therefore, routing using $k = 50$ and $k = 3$ increases the probability that packets can be routed to alternative group members in the neighborhood when the target group member becomes unreachable without compromising the traveled distance. On the other hand, routing strategies

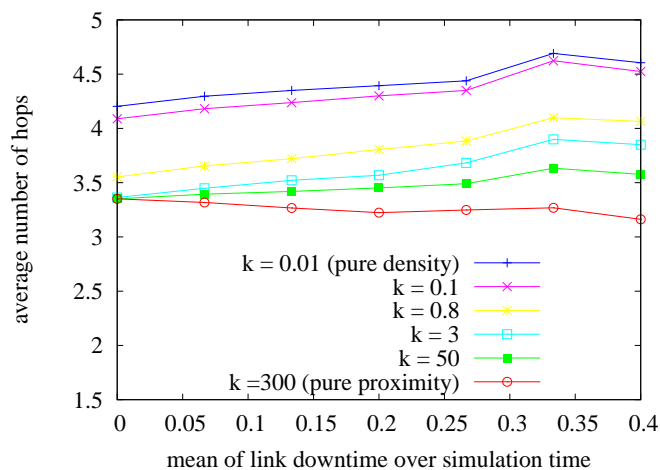
which put more weight on density such as $k = 0.1$ and $k = 0.01$ do not manage to significantly increase the packet delivery ratio because the increased probability that packets can be routed to alternative group members when the target group member becomes unreachable is compensated by the larger path lengths which increase the probability that a node along the path moves in such a way that the path becomes unavailable. Note also that for high node speeds, the traveled distance increases more for density-based routing strategies than for pure proximity-based routing, however, they deliver more data packets successfully at the same time.

The results using the random walk mobility model [17] instead of the random waypoint mobility model are shown in Figure 8. With this mobility model, each node moves a constant distance of 25 meters and then changes its direction according to an angle taken from a uniform distribution between 0 and 360 degrees. This mobility model shows a more local motion pattern than the random waypoint mobility model.

The results are similar to the simulations obtained with the random waypoint mobility model, however, with two minor differences. Due to the different spatial node distribution of



(a) Robustness.



(b) Travelled distance of delivered data packets.

Fig. 9. Robustness and path stretch for static network with temporarily link failures.

the nodes compared to the random waypoint mobility model (as observed in [15]), the traveled distance of the data packets is slightly higher on average. This in turn causes more packets to be lost for high node speeds. However, the trends are similar to the results from the random waypoint mobility model.

Finally, we consider a static scenario with temporary link failures. These conditions are typical for mesh networks where the links become temporarily unavailable due to multi-path fading effects (e.g., as reported in [19]). The simulation settings are identical to the ones described above except the nodes are at a random position in the simulation square and are not mobile. The link failures are modeled according to the two-state Markov connectivity model proposed in [18]. In essence, links are modelled as two-state random variables with an up and down state. The link lifetimes follows an exponential distribution with a mean value (the same for all links) we vary to obtain different degrees of dynamism.

The results are shown in Figure 9. Figure 9(a) plots the packet delivery ratio versus the mean of the link downtime divided by the whole simulation time. The performance degradation follows the same behavior as for the random waypoint

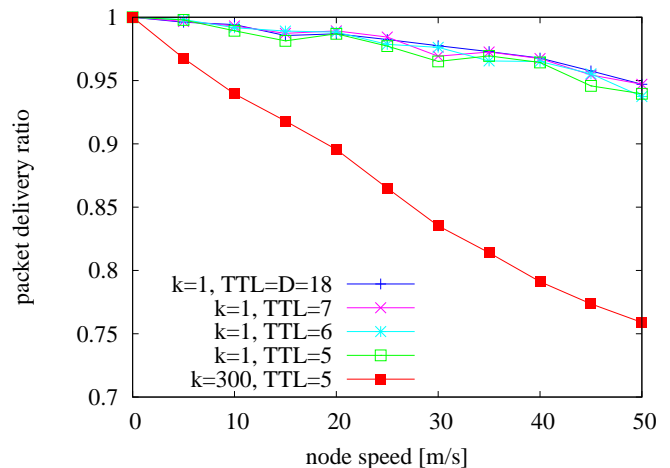


Fig. 10. The effect of different TTL values in the group member advertisement messages.

and the random walk mobility models. We conclude that density-based routing is also more robust in static networks with link failures. Also the traveled distance (see Figure 9(b)) follows the same trend as for node mobility.

To conclude, density-based routing increases the packet delivery ratio or simply the robustness. Pure density-based routing however, increases significantly the path length compared to pure proximity-based routing. The best tradeoff between robustness and path length is obtained with combined proximity- and density-based routing using a value of $k \approx 1$. In this regime, the packet delivery ratio is almost as good as pure density-based routing while keeping the path length almost as low as with proximity-based routing. This conclusion holds for our experiments with mobile networks as well as for static networks with temporarily link failures.

C. Scalability

While density-based routing improves the robustness of anycast routing, it also requires more knowledge about the available group members compared to proximity-based routing. With proximity-based routing it is sufficient to know about one member (the closest group member) whereas density requires knowledge about all group members. To assess this scalability concern, we analyze methods to limit the protocol overhead. Specifically, we examine the effect of limited flooding; i.e., when the periodic advertisement messages from the group members are no longer flooded within the entire network, but only within a limited scope. Ideally, the performance should not be significantly affected by limiting the flooding scope, and we could thus reduce the message exchange overhead.

Figure 10 compares the packet delivery ratio for $k = 1$ and different time to live (TTL) values with the packet delivery ratio of pure proximity routing ($k = 300$). The shown results are for networks of 400 nodes moving according to the random waypoint mobility model on a square of size 2500 by 2500 meters and a wireless range of 225 meters. The anycast group consists of 10 members also moving according to the random waypoint model. The network diameter is $D = 18$. We can see

that reducing the TTL value has little effect on the performance of density-based routing. Note that we do not use $TTL < 5$ because in the simulated scenarios, the shortest node-group member distance can be up to 5 hops. We conclude that it is possible to profit from the increased robustness of density-based routing without flooding the whole network and without significantly increasing the protocol overhead compared to pure proximity-based routing.

VI. RELATED WORK

The basic idea of density-based anycast routing was first proposed in an earlier version of this paper [13]. In this paper, we extend the idea and further provide a distributed field establishment protocol as well as a simulation-based evaluation of the different anycast routing strategies in three dynamic networking scenarios.

IP anycast was originally proposed in 1993 for IPv4 [1]. Later, it was incorporated into the IPv6 addressing architecture [20]. The reasons why IP anycast is not widely deployed nowadays in the Internet is mainly due to two problems. IP anycast, as originally proposed, does not scale to the number of groups and it is hard to deploy it at a large scale on the Internet infrastructure. For the first problem, Katabi et. al. proposed GIA [21] as a scalable IP architecture. Recent ideas to make IP anycast more easily deployable were proposed in [22] and [23]. Although scalability to the number of groups and ease of deployment are fundamental to the success of anycast in the Internet, it is not as critical for wireless ad hoc networks, which is where our work is targeted. The applications we have in mind, such as accessing the Internet via gateway nodes in the ad hoc network, or gathering sensor data to a few data sinks, typically require just a few groups. Furthermore, ad hoc networks do not rely on an existing routing infrastructure, which makes the problem of protocol deployment much easier.

Vincent Park et. al. described in [5] and [24] how to extend known unicast routing techniques such as link-state, distance-vector, and link-reversal routing for anycast delivery. However, by extending unicast routing techniques, it is only possible to build anycast routing protocols which find the shortest route to the closest group member.

Jianxin Wang et al. proposed in [8] and [10] to extend AODV [25] and DSR [14] respectively to support anycast delivery in mobile ad hoc networks. These protocol strategies can be instantiated with our model by choosing a large value of k . The main drawback of these protocols is again that they only consider the closest group member when computing routes.

Field-based or gradient-based routing has been proposed in the past for various type of applications including load balancing [26] in the Internet, data collection in sensor networks [27], [28], sensor node placement [29], guided navigation [30], or service discovery [31]. The novelty of our work is in the way we apply field-based routing for anycast routing. The additive effect from the different group members in our field definition results in a novel density-based routing scheme.

VII. CONCLUSIONS

In this paper, we have examined the existing anycast routing strategies and introduced a new class of anycast routing

schemes: density-based routing. We have presented a field-based routing model that represents both, the existing anycast routing schemes as well as the density-based ones. We use the results from the model evaluation to categorize the routing strategies into three regimes: (I) proximity-based routing; (II) routing as the tradeoff between proximity and density; and (III) pure density-based routing.

Our results show that density-based anycast routing is of particular interest in wireless and mobile ad hoc networks. Due to the dynamic nature of these networks, traditional proximity-based routing schemes often fail to find alternative routes when a group member moves away or when intermediate links along the path to a group member break. Under these conditions, density-based anycast routing outperforms proximity-based routing in terms of successful packet delivery because the probability to be able to re-route packets is increased when forwarding towards a high population of group members.

From our simulation studies we learn that the best routing strategy lies in a tradeoff between proximity and density, obtained using a value of $k \approx 1$ in our model. This particular tradeoff offers the increased robustness of density-based routing without introducing a significant path stretch compared to pure proximity-based routing.

It is noteworthy that many potential fields in physics such as the electric field or the gravitational field follow a potential decreasing law with a value of $k = 1$. It seems that physical laws can inspire us to design better systems and algorithms.

REFERENCES

- [1] C. Partidge, T. Mendez, and W. Milliken, "Host Anycasting Service," IETF RFC 1546, November 1993.
- [2] T. Hardie, "Distributing Authoritative Name Servers via Shared Unicast Addresses," RFC 3258, April 2002.
- [3] D. Kim, D. Meyer, H. Kilmer, and D. Farinacci, "Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)," RFC 3446, January 2003.
- [4] J. Moy, "OSPF Version 2," IETF RFC 2328, April 1998.
- [5] V. Park and J. Macker, "Anycast Routing for Mobile Services," in *Conference on Information Sciences and Systems (CISS)*, Baltimore, MD, USA, March 1999.
- [6] G. Malkin, "Rip version 2," IETF RFC 2453, November 1998.
- [7] V. Park and S. Corson, *Temporally-Ordered Routing Algorithm (TORA)*, IETF Internet Draft, July 2001.
- [8] J. Wang, Y. Zheng, and W. Jia, "An AODV-based Anycast Protocol in Mobile Ad Hoc Network," in *Proc. of the IEEE International Symposium on Personal, Indoor and Mobile Radio Communication*, Beijing, China, September 2003.
- [9] C. Intanagonwivat and D. D. Lucia, "The Sink-Based Anycast Routing Protocol for Ad Hoc Wireless Sensor Networks," USC Computer Science, CA, USA, Tech. Rep. 99-698, 1999.
- [10] J. Wang, Y. Zheng, and W. Jia, "A-DSR: A DSR-based Anycast Protocol for IPv6 Flow In Mobile Ad Hoc Networks," in *Proc. of the IEEE Vehicular Technology Conference*, Orlando, Florida, USA, October 2003.
- [11] U. C. Kozat and L. Tassiulas, "Network Layer Support for Service Discovery in Mobile Ad Hoc Networks," in *Proceedings of the IEEE INFOCOM*, San Francisco, USA, April 2003.
- [12] E. W. Dijkstra, "A Note on Two Problems in Connexion with Graphs," *Numerische Mathematik*, vol. 1, pp. 269–271, 1959.
- [13] V. Lenders, M. May, and B. Plattner, "Density-based vs. Proximity-based Anycast Routing for Mobile Networks," in *IEEE INFOCOM*, Barcelona, Spain, April 2006.
- [14] D. Johnson and D. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," in *Mobile Computing*, T. Imelinsky and H. Korth, Eds. Kluwer Academic Publishers, 1996, vol. 353, pp. 153–181.

- [15] C. Bettstetter, G. Resta, and P. 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.
- [16] 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.
- [17] T. Camp, J. Boleng, and V. Davies, "A Survey of Mobility Models for Ad Hoc Network Research," in *Wireless Communications and Mobile Computing (WCMC): Special issue on mobile ad hoc networking: research, trends and applications*, vol. 2, 2002, pp. 483–502.
- [18] T. Lin, S. F. Midkiff, and J. S. Park, "Mobility versus Link Stability in Simulation of Mobile Ad Hoc Networks," in *Proceedings of Communication Networks and Distributed Systems Modeling and Simulation Conference*, 2003, pp. 3–8.
- [19] D. Aguayo, J. Bricket, S. Biswas, G. Judd, and R. Morris, "Link-level Measurements from an 802.11b Mesh Network," in *Proceedings of ACM SIGCOMM'04*, Portland, Oregon, USA, October 2004.
- [20] S. Deering and R. Hinden, "IP Version 6 Addressing Architecture," IETF RFC 2373, July 1998.
- [21] D. Katabi and J. Wroclawski, "A Framework for Scalable Global IP-Anycast (GIA)," in *Proc. of ACM SIGCOMM*, Stockholm, Sweden, August 2000.
- [22] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, "Internet Indirection Infrastructure," in *Proceedings of ACM SIGCOMM Conference*, Pittsburgh, PA, USA, August 2002.
- [23] H. Ballani and P. Francis, "Towards a Global IP Anycast Service," in *Proc. of ACM SIGCOMM*, Philadelphia, USA, August 2005.
- [24] V. D. Park and J. P. Macker, "Anycast Routing for Mobile Networking," in *Proc. of MILCOM*, Atlantic City, New Jersey, USA, November 1999.
- [25] C. E. Perkins, E. M. Belding-Royer, and S. R. Das, "Ad Hoc On-Demand Distance Vector (AODV) Routing," IETF Internet Draft, draft-ietf-manet-aodv-12.txt, November 2002.
- [26] A. Basu, A. Lin, and S. 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.
- [27] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," in *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (Mobicom '00)*, Boston, USA, 2000.
- [28] J. Faruque and A. 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.
- [29] S. Toupis and L. Tassiulas, "Packetostatics: Deployment of Massively Dense Sensor Networks as an Electrostatic Problem," in *IEEE INFOCOM*, Miami, USA, March 2005.
- [30] 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.
- [31] V. Lenders, M. May, and B. Plattner, "Service Discovery in Mobile Ad Hoc Networks: A Field Theoretic Approach," *Elsevier Journal on Pervasive and Mobile Computing*, vol. 1, no. 3, September 2005.



Martin May (M'97) received the Master degree in computer science from the University of Mannheim in 1996. In 1999, he received his Ph.D. degree from the University of Nice, France. He did most of his thesis work on Internet QoS mechanisms at INRIA, Sophia Antipolis, France, but was also technical staff member of Lucent Bell-Labs Research, Holmdel, USA and Sprintlabs, Burlingame, USA.

Until beginning of 2000, he continued his research as a post-doctoral member of the research staff at Sprintlabs, Burlingame, US. From 2000 until 2003, he founded a start up company in France where he worked in the field of Content Delivery Networking. He is now senior research assistant at the Swiss Institute of Technology in Zurich (ETHZ). His research interests are in self-organizing networks, mobile ad hoc networks and network security.

Dr. May is a member of the IEEE, ACM and the Internet Society. He served on the program committees for multiple networking conferences and is an active member of the IETF where he is co-editor of multiple Internet-drafts and RFCs.



Bernhard Plattner (S'75-M'83) is a Professor of Computer Engineering at ETH Zurich, where he leads the communication systems research group. His research currently focuses on self-organizing networks and systems-oriented aspects of information security. He has been the principal investigator or Co-PI of numerous national and international projects.

In 1996-98, Dr. Plattner served as the Head of Faculty of Electrical Engineering at ETH Zurich. He currently is the Vice-Rector for Bachelor/Master

Studies at ETH Zurich.

Dr. Plattner is a member of the IEEE, ACM and the Internet Society. He served as the program or general chair of various international conferences, such as ACM SIGCOMM 91, INET '94, and IWAN 2002.



Vincent Lenders received his Masters in electrical engineering in 2001 from ETH Zurich, Switzerland and his Ph.D. in 2006 also from ETH Zurich.

After his Ph.D. degree, he joined Princeton University, New Jersey, USA where is currently working as a post-doctoral research fellow. His main research interests include mobile, ad hoc, and delay-tolerant networks.

Dr. Lenders served on the program committees for multiple IEEE and ACM networking conferences.