



Available at

www.ElsevierComputerScience.com

POWERED BY SCIENCE @ DIRECT®

Ad Hoc Networks 1 (2003) 351–369

Ad Hoc
Networks

www.elsevier.com/locate/adhoc

Contention-based forwarding for mobile ad hoc networks [☆]

Holger Füßler ^{a,*}, Jörg Widmer ^a, Michael Käsemann ^a, Martin Mauve ^a,
Hannes Hartenstein ^b

^a Department of Mathematics and Computer Science, University of Mannheim, L15, 16, D-68161 Mannheim, Germany

^b NEC Network Labs Europe, Kurfürsten-Anlage 36, D-69115 Heidelberg, Germany

Abstract

Existing position-based unicast routing algorithms which forward packets in the geographic direction of the destination require that the forwarding node knows the positions of all neighbors in its transmission range. This information on direct neighbors is gained by observing beacon messages each node sends out periodically.

Due to mobility, the information that a node receives about its neighbors becomes outdated, leading either to a significant decrease in the packet delivery rate or to a steep increase in load on the wireless channel as node mobility increases. In this paper, we propose a mechanism to perform position-based unicast forwarding without the help of beacons. In our *contention-based forwarding scheme* (CBF) the next hop is selected through a distributed contention process based on the actual positions of all current neighbors. For the contention process, CBF makes use of biased timers. To avoid packet duplication, the first node that is selected suppresses the selection of further nodes. We propose three suppression strategies which vary with respect to forwarding efficiency and suppression characteristics. We analyze the behavior of CBF with all three suppression strategies and compare it to an existing greedy position-based routing approach by means of simulation with ns-2. Our results show that CBF significantly reduces the load on the wireless channel required to achieve a specific delivery rate compared to the load a beacon-based greedy forwarding strategy generates.

© 2003 Elsevier B.V. All rights reserved.

Keywords: Mobile ad hoc networks; Position-based routing; Beaconless forwarding

1. Introduction

The idea of position-based routing was originally developed for packet radio networks in the 1980s [1,2]. Due to the availability of GPS it received renewed interest during the last few years as

a method for routing in mobile ad hoc networks [3–6]. The general idea of position-based routing is to select the next hop based on position information such that the packet is forwarded in the geographical direction of the destination.

The most important characteristic of position-based routing is that forwarding decisions are based on local knowledge. It is not necessary to create and maintain a global route from the sender to the destination. Therefore, position-based routing is commonly regarded as highly scalable and very robust against frequent topological changes. It is particular well suited in environments

[☆] This work was supported by the 'FleetNet' project as part of German Ministry of Education and Research (BMBF) contract no. 01AK025D.

* Corresponding author.

E-mail address: fuessler@informatik.uni-mannheim.de (H. Füßler).

where the nodes have access to their geographical position, such as in inter-vehicle communication [7,8].

Position-based routing can be divided into two main functional elements: a *location service* and a *position-based forwarding strategy*. The location service maps the unique identifier (such as an IP address) of a node to its current geographical position. It can be seen as analogous to the route discovery process of reactive topological routing algorithms such as DSR [9] or AODV [10]. For the remainder of this work we assume that an appropriate location service is present which supplies the sender of a packet with the geographical position of the packets' destination. Candidates for location services are outlined in the section on related work.

Position-based forwarding is performed by a node to select one of its neighbors in transmission range as the next hop the packet should be forwarded to. Usually, the forwarding decision is based on the node's own geographical position, the position of all neighbors within transmission range and the geographical position of the destination. The sender requests the position of the destination from the location service and then includes it in the header of the packet. Given this information, the node forwards the packet to one of its neighbors such that the packet makes progress toward the destination. This process is called *greedy forwarding*. It is possible that there is no neighbor with positive progress toward the destination while a valid route to the destination exists. The packet is then said to have reached a local optimum. In this case, a *recovery strategy* is used to escape the local optimum and to find a path toward the destination.

In all existing strategies for greedy unicast forwarding, the position of a node is made available to its direct neighbors (i.e., nodes within single-hop transmission range) in form of periodically transmitted beacons.¹ Each node stores the information it receives about its neighbors in a table and

thus maintains position information about all direct neighbors.

While the beaconing frequency can be adapted to the degree of mobility the fundamental problem of inaccurate position information is always present: a neighbor selected as a next hop may no longer be in transmission range. As will be outlined later (see Section 4, Fig. 12) this leads to a significant decrease in the packet delivery rate with increasing node mobility and to a high load on the wireless channel due to several MAC layer retransmissions. To reduce the inaccuracy of position information it is possible to increase the beaconing frequency. However, this also increases the load on the network up to a point where the available capacity is almost exclusively used for the transmission of beacons. Alternatively, it has been proposed to hand packets back to the routing layer if the next hop is no longer available [3]. At the routing layer the packets are then rerouted to a different neighbor. While this eliminates the problem of packet drops, the trial-and-error approach can cause even more bandwidth-consuming MAC layer retransmissions. Our experiments (see Fig. 14) indicate that under high mobility, the beacon-based forwarding approach requires on average more than three MAC transmissions for one single-hop packet forwarding, increasing the load on the network caused by data packets by more than a factor of three. Existing work (e.g., [3]) does not take this effect into account since there the load is measured at the routing level instead of the MAC layer.

Thus, for a given packet delivery rate, the load at the MAC layer increases dramatically with beacon-based greedy unicast forwarding (either through an increased beaconing frequency or through trial-and-error) with increasing node-mobility. In addition, a node forwarding a packet can only select a neighbor as next hop if it is contained in its neighbor table. Nodes that just moved into transmission range and that have not yet sent a beacon are therefore not considered as next hop nodes. This may lead to the failure of greedy forwarding even though an appropriate neighbor is present.

In this paper, we propose a novel greedy forwarding strategy for position-based routing algo-

¹ There are position-assisted approaches that do not require beacons (e.g., LAR) but they do require directional flooding and can thus not be considered to be *unicast* forwarding.

rithms. We call the approach *contention-based forwarding* (CBF). CBF performs greedy forwarding without the help of beacons and without the maintenance of information about the direct neighbors of a node. Instead, all suitable neighbors of the forwarding node participate in the next hop selection process and the forwarding decision is based on the actual position of the nodes at the time a packet is forwarded. This is in contrast to existing greedy forwarding algorithms that base their decision on the positions of the neighbors as they are perceived by the forwarding node. In order to escape from local optima, existing recovery strategies, as mentioned in the section on related work, can either be used directly or may be adapted to be used with CBF.

CBF shows advantages over existing greedy forwarding strategies in two important aspects:

1. *Use of accurate position information:* in CBF each neighbor uses the (very accurate) position information it has about itself to determine if it should become the next hop for a given packet. For a given delivery rate the required bandwidth for CBF does therefore not increase with node mobility (i.e., neither an increased beaconing frequency, nor trial-and-error is needed). In addition, CBF always bases the selection of the next hop on *all* direct neighbors, even those that have just moved into transmission range.
2. *Elimination of beacon overhead:* removing the beacons eliminates a major part of routing overhead which occurs independently of the actual data traffic. This includes the bandwidth used for the transmission of beacons² and the memory required in the nodes to store neighbor information.

CBF consists of two parts: the *selection* of the next hop is performed by means of contention, while *suppression* is used to reduce the chance of accidentally selecting more than one node as the

next hop. We present three suppression strategies with different suppression characteristics. The results of our study show that suppression of duplicate packets works well, that CBF has similar packet delivery ratios as beacon-based greedy routing, and that it dramatically reduces the load on the wireless medium for a given delivery rate if node mobility is high. CBF, therefore, represents a good alternative to traditional beacon-based greedy forwarding.

The contention process of CBF used for next-hop selection represents a paradigm change in the forwarding of packets. In traditional protocols, the forwarder actively selects the desired next-hop by unicasting the packet to the corresponding MAC address. In contrast, with CBF the responsibility for next-hop selection lies with the set of possible next hops. Furthermore, if no other interaction between forwarder and next hop is required, which is the case for two of the three presented strategies, MAC layer addresses become obsolete.

The remainder of this paper is structured as follows: In Section 2 we summarize related work. Section 3 contains a description of CBF with three alternative suppression schemes. In Section 4, the properties of CBF are analyzed and its performance is further investigated in Section 5 by means of simulation. Finally, Section 6 points out directions of future work and concludes the paper.

2. Related work

Ad hoc routing protocols can be classified into topology-based [11] and position-based schemes [12]. In addition, they can be further subdivided into proactive and reactive methods: with a proactive method, routing information is maintained independently of actual data communication. With a reactive scheme routing information is acquired on-demand when there is data to forward.

As described in the introduction, position-based routing consists of the forwarding algorithm and a location service which maps the ID of a node to its geographic position. The forwarding algorithm can be further subdivided into greedy forwarding and a repair strategy for situations where greedy

² While some existing MAC protocols do require beacon messages (e.g., for synchronization purposes), the overhead incurred by these beacons is very small compared to that required for beacon messages used for building up neighbor tables.

forwarding fails. Examples for position-based forwarding algorithms are face-2 [6], Greedy perimeter stateless routing (GPSR) [3], and Terminodes routing [5]. Existing position-based greedy forwarding algorithms have both proactive and reactive elements: they require proactive information about the local neighborhood of a node. This information is provided in form of periodic beacon messages transmitted independently of data packets. The actual position-based greedy forwarding of data packets is then done in a reactive manner. The CBF algorithm put forward in this paper does not require beacons and thus completely eliminates the proactive part of position-based routing. It is an alternative to the greedy forwarding part of existing protocols. An approach similar to area-based suppression, one of the three suppression schemes presented here, was independently proposed in [13]. As a repair strategy Geodesic Anchors as proposed for Terminodes routing could be used. The repair strategies of face-2 and GPSR would require a heuristic for traversing graphs without the need to know a node's neighbors.

Another related routing algorithm is location aided routing (LAR) [14]. LAR is a reactive topology-based routing algorithm, employing position information only to limit network load during the route discovery phase. A route request in LAR is flooded in the direction of the destination. This directed flooding does not require beacons: when a node receives a route request it checks whether it is in the region that leads to the destination. If this is the case it forwards the request. While directional flooding is a robust approach for route discovery, it does not scale well being used for unicast transmissions. The key difference between LAR and CBF is that CBF performs suppression to avoid packet duplication and to provide unicast capability, while the route request scheme of LAR uses directional flooding where packet duplication is common and desired.

Homezone [15], the grid location service (GLS) [16], and the location service part of DREAM [4] are examples for existing location services. Although some location services (like GLS) distribute location information in a proactive fashion, one can think of mechanisms that are fully reac-

tive, for example a scheme similar to a DSR route request [9], where a node issues a location request using flooding and the node with the requested ID answers with its current position. A reactive location service designed along these lines was analyzed in [17]. In combination with our forwarding scheme a completely reactive position-based routing algorithm can be designed.

3. Contention-based forwarding

The general idea of CBF is to base the forwarding decision on the current neighborhood as it exists in reality and not as perceived by the forwarding node. This requires that all suitable neighbors of the forwarding node are involved in the selection of the next hop.

CBF works in three steps: first, the forwarding node transmits the packet as a single-hop broadcast to all neighbors.³ Second, the neighbors compete with each other for the “right” to forward the packet. During this *contention period*, a node determines how well it is suited as a next hop for the packet. Third, the node that wins the contention *suppresses* the other nodes and thus establishes itself as the next forwarding node.

In the following we describe in detail how contention can be realized on the basis of biased timers. Furthermore, we present three different suppression strategies.

3.1. Timer-based contention

The decentralized selection of one node out of a set of nodes is a common problem encountered in many areas of computer networks. It is known as feedback control in group communication [18, 19] or as medium access control in (wireless and

³ In general, this should require similar resources as a single-hop unicast transmission except that packets for other nodes cannot be discarded at the network interface but have to be passed up the protocol stack. Depending on the physical and MAC layer there may be further differences between unicast and broadcast (e.g., in IEEE 802.11 the sleep mode may not be applicable).

wired) local area networks such as IEEE 802.11 [20].

A standard approach for this selection is by means of timers. In its most simple form, timer-based contention requires that each node sets a timer with a random value. Once the first timer expires, the corresponding node responds. The timers of all other nodes are canceled and their responses are suppressed.

It is important to realize that with this contention algorithm more than one node may respond, even if a ‘good’ suppression mechanism is used. This will happen when the difference between the timeout value of the earliest timer and some other timer is smaller than the time required for suppression. Therefore, the interval from which the timeout values are selected should increase with the number of competing nodes. It was shown in [18] that exponentially distributed random timers can further decrease the number of responses compared to uniformly distributed timers.

To use such a simple timer-based mechanism for the forwarding decision, all nodes that receive the packet check if they are closer to the destination than the forwarding node. If this is the case, a random (exponentially distributed) timer is set to start the contention and the node responding first is selected as the next hop.

The problem of the simple timer-based contention is that all nodes which are located closer to the destination than the forwarding node are treated equally. Thus a node providing minimal progress would have the same chance to be selected as next hop as a node providing a large progress. We therefore propose to determine the value for the timers based on how much progress a node provides toward the destination instead of setting them randomly.

To greedily minimize the remaining distance to the destination, the progress P is defined as ⁴

$$P(f, z, n) = \max \left\{ 0, \frac{\text{dist}(f, z) - \text{dist}(n, z)}{r_{\text{radio}}} \right\}$$

⁴ Note that the original definition of progress in [2] is different to ours since in [2] an additional projection onto the line crossing f and z is used.

given f is the position of the forwarder, z the position of the destination and n the position of the considered neighbor. dist is defined as the Euclidean distance between two positions and r_{radio} is the nominal radio range.

Fig. 1 illustrates the suitability of a node as next hop depending on its location. A progress value (P) of 0 indicates that a node is unsuitable while a value of 1 is optimal and is reached if the node is located at the intersection of the circle delineating the transmission range of the forwarding node and the line from the forwarding node to the destination. Thus P increases linearly from 0 to 1 with the progress that a node at this position would provide for the packet.

For the contention in CBF we select the timer runtime as

$$t(P) = T(1 - P),$$

where T is the maximum forwarding delay. This makes sure that the node with the largest progress is selected as next hop. Since the runtime of the timer only depends on the remaining distance to the destination it is identical for all nodes that are located on the same circle around the destination. A packet duplication may occur in the following situation: if the best suited node has a progress of P_1 and there exists at least one node with a progress of P such that $t(P) - t(P_1) < \delta$, where δ is the minimum time interval needed for suppression, then at least one packet duplication occurs. All nodes with progress P and

$$P_1 \geq P \geq 1 - \frac{\delta + T(1 - P_1)}{T} = P_1 - \frac{\delta}{T}$$




Fig. 1. Packet progress (transmission range 250 m).




Fig. 2. Duplication area.

are within this so-called *duplication area* and cannot be suppressed, as shown in Fig. 2.

An interesting property of the duplication area is that it becomes smaller the closer the best suited node is located to the destination. As long as the positions of the nodes are uniformly distributed this reduces the chance of packet duplication in a similar way as exponentially distributed random timers reduce the chance of packet duplication when compared to linearly distributed random timers.

Analytically, this property can be made explicit via the probability density function (PDF) of the progress of a randomly selected point within the forwarding node's transmission range. Let d denote the distance between forwarding node and destination and let us assume a normalized transmission range of 1. The radius r of a circle around the destination as depicted in Fig. 2 corresponds to a progress $d - r$ for $r \in [d - 1, d + 1]$. The PDF for progress $d - r$ is given by

$$\frac{2}{\pi} \cdot r \cdot \arccos\left(\frac{r^2 + d^2 - 1}{2dr}\right). \quad (1)$$

Graphs of expression (1) for $d = 1, 2, 20$ are shown in Fig. 3.⁵ From the shape of these graphs it can be seen that there are relatively few well suited nodes (with a large positive progress). Setting the contention timer according to the progress will

⁵ We note that this figure ignores that values below zero are unsuitable for forwarding.




Fig. 3. PDF of packet progress.

thus result in few timers with a short runtime and many timers with a long runtime which decreases the likeliness of packet duplication.

Packet duplication is closely coupled with the characteristics of the MAC layer. With many MAC schemes (as for example IEEE 802.11), packets will be serialized and thus packet duplication can be avoided. In wireless networks based on CSMA/CA [21], the serialization is not only performed between packets from nodes which are in transmission range of each other but is typically done on the basis of the interference range which is roughly twice the transmission range. As a consequence, the transmission of all neighbors of the forwarding node will be serialized since the distance between any two neighbors does not exceed twice the transmission range. If packets can be removed from the interface queue of the MAC layer, then the forced serialization can be used to eliminate the effect of packet duplication caused by the suppression delay δ , as described in Section 3.1. One node will be the first to forward a packet. Other nodes that have queued a duplicate of the packet may drop it once they overhear the forwarding of the packet by another node.

3.2. Suppression

Let us now assume that all neighbors of the forwarding node have set their contention timer according to their respective distances to the des-

mination. After the first of those timers expires, a suppression algorithm aims to cancel the timers in all other nodes to prevent multiple next hops and thereby packet duplication.

3.2.1. Basic suppression scheme

The most basic conceivable suppression mechanism works as follows: if the timer at a node expires, the node assumes that it is the next hop and broadcasts the packet. When another node receives this broadcast and still has a timer running for the packet, the timer is canceled and the node will not forward the packet.

Depending on where the initial next hop is located, other nodes may be out of transmission range and will thus not be suppressed. In the worst case, up to three copies of the packets may be forwarded, as shown in Fig. 4. The larger the number of nodes within transmission range of the source, the higher the probability of one or more packet duplications.

It should be noted that the packet duplications described here are in addition to packet duplications caused by the amount of time required for the suppression of other nodes, as described in the previous section. They do occur even if the suppression requires no time at all.

3.2.2. Area-based suppression

In order to avoid the extra packet duplications from the basic suppression scheme we propose to artificially reduce the area from which the next hop is selected. We call this reduced area the *suppression*




Fig. 4. Packet duplication in the basic scheme.




Fig. 5. Forwarding areas.

area and the algorithm *area-based suppression*. The key idea is to choose the suppression area such that all nodes within that area are in transmission range of each other, avoiding extra packet duplications as they may appear in the basic suppression scheme.

Area-based suppression requires a decision on how the suppression area is chosen. One possible choice is a circle with the diameter of the transmission range located within the forwarding node's transmission range in direction of the destination (e.g., the gray circle in Fig. 5). A circle is the geometric shape covering the largest area given that any two points within the shape are no farther apart than the transmission range. If the nodes are uniformly distributed this means that on average the circle will contain the highest number of neighboring nodes when compared to other shapes where the distance between any two points does not exceed the transmission range. However, several parts of the forwarding area which make good forwarding progress are not included in the circle. A different shape where any two points are no farther apart than the transmission range, the Reuleaux triangle [22], much better covers the area with good forwarding progress (see Fig. 5).⁶ By using the Reuleaux triangle with a width of the transmission range, we trade off the number of nodes contained in the suppression area against the inclusion of better suited nodes. The motivation

⁶ A Reuleaux triangle with a width of r can be constructed by placing three circles with radius r at the corners of an equilateral triangle with an edge length r . The intersection of the circles is the Reuleaux triangle.




Fig. 6. PDF of nodes with equal forward progress (total) and fractions contained within the circle and Reuleaux areas.

for using the Reuleaux triangle is illustrated in Fig. 6. The curve titled “total” is the PDF for the progress of nodes with positive progress. The curve “circle” denotes the fraction of the density “total” for a neighbor with progress p to be contained in the circle. The same applies for the “Reuleaux” curve and the Reuleaux triangle. Between 60% and 100% progress, the Reuleaux triangle covers more of the neighbors than the circle and above approximately 80%, the Reuleaux triangle covers all of the neighbors with this progress. Therefore, it is more likely to include a node with good forwarding progress.

Given the Reuleaux triangle as suppression area, the suppression algorithm works as follows:

- The forwarding node broadcasts the packet.
- Only the nodes contained in the Reuleaux triangle participate in the contention process.
- The node at which the timer runs out first is the next hop and broadcasts the packet.
- All other nodes are suppressed. Packet duplication may occur only because of the time required for suppression.

Of course it is possible that the only neighbors of the forwarding node that provide forward progress toward the destination are not contained in the Reuleaux triangle (1). In this case the forwarding node will not hear another node forwarding the packet. Consequently, the process is

repeated with the remaining areas (2) and (3) where nodes with forwarding progress may be located, until the forwarding node hears a rebroadcast of the packet. If no node within areas (1), (2), or (3) responds, then there is no node with positive forward progress and a recovery strategy has to be used just like in existing position-based forwarding schemes. The order in which areas (2) and (3) are selected when no node is located in area (1) should be chosen randomly. This way, a tendency to always route around areas with little or no coverage in the same direction is avoided.

The key advantage of area-based suppression is the reduction of packet duplications. This comes at the cost of requiring up to three broadcasts for forwarding a packet. However, it is important to realize that requiring more than one broadcast becomes less and less likely as the number of nodes increases. Furthermore, the Reuleaux triangle covers the largest of the three areas and therefore has the highest probability of containing a potential next hop.

3.2.3. Active selection

While area-based suppression eliminates the packet duplications caused by nodes not being in transmission range of each other it does not prevent packet duplications caused by the time required to perform the suppression. Active selection of the next hop prevents all forms of packet duplication at the cost of additional control messages. It is inspired by the request to send, clear to send (RTS/CTS) MACA-scheme proposed in [23] and used (as a variant) in IEEE 802.11 (see [20]).

The scheme works as follows: the forwarding node broadcasts a control packet called RTF (request to forward) instead of immediately broadcasting the packet. The RTF contains the forwarding node’s location and the final destination’s location. Every neighbor checks if it provides forward progress for the packet announced by the RTF. If this is the case it sets a reply timer according to the basic suppression scheme. If the timer runs out, a control-packet called CTF (clear to forward) is transmitted to the forwarding node. The CTF packet contains the position of the node sending the CTF. If a node hears a CTF for the packet, it deletes its own timer and is suppressed.

The forwarding node may receive multiple CTF control-packets. Of all neighbors that have transmitted a CTF packet it selects the node with the largest forward progress and transmits the packet to this node using unicast. An additional benefit of active selection compared to basic and area-based suppression is that it may be integrated with RTS/CTS schemes to avoid the “hidden terminal problem”.

Active selection prevents all forms of packet duplication, even though multiple nodes may send a CTF control packet. The forwarding node acts as a central authority deciding which node is selected as the next hop. This comes at the cost of additional overhead in form of RTF/CTF control packets.

4. Performance analysis

A very most important characteristic of the different algorithms is the packet duplication probability. Furthermore, it is interesting to see how much message overhead and time is required to forward a packet from hop to hop. In the following we determine the likeliness of packet duplication and the forwarding delay for each of the three suppression schemes.

For the analysis, the following model was used. Without loss of generality, the forwarding node is located at position (0,0) and the transmission range is set to one. The position of the final destination is (d_x, d_y) with d_x and d_y uniformly distributed in $[0;20)$ and $1 \leq \sqrt{d_x^2 + d_y^2} < 20$. Neighbor nodes are sampled similar with the number of neighbors increasing exponentially from 1 to 256. The timer used for contention is calculated by each neighbor n with forward progress as

$$t(n) = T \left(1 - \frac{\sqrt{(d_x - n_x)^2 + (d_y - n_y)^2}}{\sqrt{d_x^2 + d_y^2}} \right), \quad (2)$$

where T is the maximum response time and $t(n) \in [0; T]$.⁷

⁷ For a reasonably low variance, each simulation was run 10^7 times. As pseudo-random number generator, the “Mersenne Twister” [24] as implemented in the GNU scientific library [25], was used.

4.1. Average number of next hops

The simulation results regarding the probability of packet duplication for the three algorithms are presented in Fig. 7. In the simulations, there is no suppression delay ($\delta = 0$) and no node mobility.

For the basic suppression scheme, there are at most three next hops and packet duplication can only occur because nodes are further apart than the transmission range and thus do not suppress each other. With a growing number of neighbors, the probability of “no next hop” approaches zero while the likeliness of packet duplication (2 or 3 next hops) increases. The probability of a single next hop with the basic suppression scheme reaches a maximum for approximately 4 neighboring nodes. With more than 9 neighbors, packets are duplicated with a probability of more than 0.5.

In area-based suppression, packet duplication can only occur due to suppression delay or node mobility and we should see no duplication otherwise. This is confirmed by the simulation results presented in Fig. 7(b). The curve for “no neighbors with forward progress” quickly drops to zero as the number of neighbors increases and in most cases exactly one node will forward the packet.

For the active selection scheme there can be no packet duplication at all, since the forwarder is the final arbiter for the decision which neighbor is selected as the next hop. This comes at the cost of additional overhead. The overhead consists of one RTF control packet transmitted by the forwarder and of one or more CTF control packets transmitted by the neighbors. The number of CTF control packets generated is the same as the number of unsuppressed nodes in the basic suppression scheme and can thus be seen in Fig. 7(a).

4.2. Impact of the suppression delay δ

For the basic and the area-based scheme, packet duplication can occur even if neighbors are within each others transmission range, as long as they are contained in the duplication area. The size of the duplication area depends on the time required for the suppression, causing an increase in packet duplication probability with increasing




Fig. 7. Relative probabilities of n next hops ($\delta = 0$). (a) Basic suppression scheme. (b) Area-based suppression and active selection.




Fig. 8. Average number of next hops for increasing suppression delay δ . (a) Basic suppression scheme. (b) Area-based suppression.

suppression delay. In Fig. 8(a), the average number of next hops for different suppression delays is shown for the basic scheme. While a suppression delay of $0.001T$ and $0.01T$ affects the duplication of packets only marginally, a suppression delay of $0.1T$ causes significant packet duplication even for low numbers of neighbors. Hence, given a certain (MAC dependent) suppression delay, T should be chosen as a large multiple of δ if the basic suppression scheme is used.

The number of duplicates is much lower when area-based suppression is used. Also, there is no significant increase in the number of next hops as long as δ is a small fraction of T . Only for $\delta = 0.1T$ there is a noticeable increase in duplicate packet as shown in Fig. 8(b).

As discussed before, active selection will not cause packet duplication due to the suppression delay.

4.3. Forwarding delay

With respect to delay, the basic suppression scheme is faster than the other two alternatives. The only delay introduced is caused by waiting for the first neighbor to forward the packet, as depicted in Fig. 9.

With the area-based suppression scheme it is possible that no node with forward progress is contained in the Reuleaux triangle oriented toward the destination, even though a neighbor with forward progress exists outside of this area. Up to




Fig. 9. Average time before next forward.

three broadcast transmissions of the same packet may be necessary to guarantee that a suitable neighbor is found if one exists. Fig. 10 shows the probability distribution for the number of broadcasts required to find a neighbor with forward progress. Again, it is possible that no neighbor with forward progress exists. From Fig. 10 we observe that for any significant number of neighbors, it is highly likely that a node is located within the Reuleaux triangle. This corresponds to the conclusions made in Section 3.2.2 concerning Fig. 6. In particular, the best nodes are likely to be located within the Reuleaux triangle.

The area-based suppression has the same characteristics as the basic suppression scheme when a forwarding node can be found in the first Reuleaux




Fig. 10. Relative probabilities of “first next hop is in region”.

triangle. Otherwise, the forwarding node has to wait for T and then has to rebroadcast the packet in the second and possibly even the third area. The probability of no next hop in the Reuleaux triangle is very small for a reasonable number of neighbors (6 or more). Hence, the difference in forwarding delay between the basic and the area-based suppression scheme is only significant for a small number of neighbors within transmission range.

The forwarding delay introduced by active selection depends not only the time required to transmit a data packet but also on the time to transmit the RTF and CTF. Both packets are likely to be small and the time to transmit them should be significantly smaller than the time for data packet transmission. If the forwarder waits for the feedback delay T (i.e., until all possible CTFs have arrived) and then forwards the packet to the best suitable node, we have a constant forwarding delay of T , as shown in Fig. 9.

With an integration of MAC layer and CBF, the forwarding delay can be improved by giving a higher priority to data packets which suppress subsequent CTF packets after the first CTF was received by the forwarder.

4.4. General remarks

To conclude, even though the basic suppression scheme is the fastest and does not incur any additional overhead in terms of additional messages or retries until a next hop is found, its applicability is limited. Even under favorable conditions packet duplication occurs with a likeliness of more than 50% at each hop. Therefore, more sophisticated suppression schemes are desirable.

The area-based suppression scheme is very well suited if the density of nodes is sufficiently high. Only for very small numbers of neighbors the good suppression characteristics are offset by a larger forwarding delay.

Active selection can be used with all node densities and suppression delay values. There will be no uncontrolled duplication of packets. Its main drawback is that it transmits at least two additional packets (RTF/CTF) for each forwarding of the data packet. For scenarios where the density of nodes is high and the suppression delay

is comparatively low the area-based suppression scheme may be preferable.

5. Protocol simulations

5.1. Simulation setup

The proposed mechanisms were implemented in the ns-2 network simulator [26] version 2.1b8a (using the MAC layer of the version 2.1b9 with additional bug fixes). The size of the simulated area is 2 km×2 km. We simulate different node densities and different levels of mobility using the Random Waypoint Model [9].⁸ The different levels of mobility are achieved by modifying the maximum node speed, with a movement pause time of zero. For every combination of protocol-variant, node density, and maximum speed we generate 50 independent sets of movement scenarios. For each of these scenarios, we randomly pick one sender-receiver pair. The sender transmits 100 packets with a payload of 128 bytes with a constant rate of 4 packets per second. Each simulation lasts for 40 s of simulation time. Data traffic starts at 5–10 s (randomized) after the start of the simulation, giving the beacon-based protocols time to exchange neighbor information and leaving enough time to deliver outstanding packets at the end before the simulation is terminated.

The simulated protocols are the three CBF schemes as described in Section 3.2 and a basic greedy forwarding mechanism based on GPSR [3]. The protocols are simulated without perimeter mode (i.e., without repair strategy if greedy forwarding fails to find a route to the destination). Greedy forwarding using beacons is simulated with and without the ability to re-route packets if a selected next-hop is not reachable by the link layer,

the so-called MAC callback option. The two alternatives are called ‘optimized greedy’ and ‘basic greedy’ in the discussion of the simulations. The simulated beacon intervals are 0.5, 1.0, and 2.0 s and both greedy schemes used implicit beaconing, where beacons are also piggybacked on data packets.

An ‘optimization’ used for CBF is the introduction of a duplication suppression scheme that works as follows: Every packet is marked with a packet ID by the original sender. If a node has already forwarded a packet with this ID or was suppressed during the contention, it will not attempt to forward the packet again.

The underlying MAC protocol is IEEE 802.11 with a capacity of 1 MBit/s to ensure that the broadcasts for CBF (as well as the beacons) and the unicast packets for greedy routing are transmitted at the same bitrate.⁹ For the simulations, an unmodified MAC 802.11 was used but we note that with an integration of CBF and MAC, the performance of CBF can be improved considerably.

In the following sections we investigate the performance of the different routing algorithms with particular focus on the impact of node mobility. A more extensive simulation study of CBF can be found in [28].

5.2. Node density

As a first sanity check, we simulated CBF and greedy forwarding without node mobility for different node densities. Without mobility, the beacon interval has no impact on the performance of greedy routing and location information is always accurate.

An immediate result of this simulation is that the runs with 100 and 200 nodes result in high packet loss rates for all approaches. This is caused by the fact they frequently reach a local optimum and thus fail with low node densities.

⁸ Note that with the random waypoint model, the node density is not uniform [27]. The higher the node mobility, the earlier will the originally uniformly distributed nodes accumulate in the middle of the simulation area, decreasing the average communication distance. Nevertheless, we choose the model to allow comparison of our simulation results with other simulation studies.

⁹ Earlier versions of the ns-2 MAC had a bug using a higher rate for broadcasts than the standard allows. This bug is fixed in the code we used.

In more detail it can be observed that the basic CBF scheme achieves a higher packet delivery ratio than all other schemes as shown in Fig. 11. Due to packet duplication, packets may be forwarded along a non-greedy path and find a route to the destination even if no greedy route exists. As is to be expected, the other CBF schemes as well as greedy forwarding have very similar packet delivery ratios, which depend mostly on the probability that a greedy route exists given the current node density. The area based scheme has a slightly lower packet delivery ratio for very low node densities as the sequence of probing areas may result in choosing a forwarding node that makes less progress than the best node of all forwarding areas. For higher node densities where the forwarder is almost always in the first forwarding area this discrepancy vanishes. Active selection performs slightly worse than the other schemes for higher node densities since the request response procedure increases the likelihood that a packet collision occurs during the forwarding process. Its performance could easily be improved by allowing packet retransmissions.

The analysis of other performance measures (e.g., routing overhead and forwarding delay) is of little value if only a fraction of the sent packets arrive at the destination. For this reason, we limit the remainder of our analysis to simulations with 300 nodes.




Fig. 11. Packet delivery ratio for different node densities.

5.3. Packet delivery ratio

Fig. 12 shows the packet delivery ratio of the three CBF schemes: the basic greedy scheme for all three simulated beacon intervals and the optimized greedy scheme for a beacon interval of one second. The values for optimized greedy with other beacon intervals were omitted because their performance in the chosen scenarios is similar to the run with a beacon interval of 1 s. The node density is 300 nodes in the simulated area of 4 km². The x-axis shows the four different groups of movement scenarios with their respective maximum node speed.

As can be seen from the graph, all CBF schemes and the optimized greedy scheme reach very high packet delivery ratios. Since the node density is fairly high, greedy routes exist most of the time. Only the packet delivery ratio of the active selection scheme suffers slightly when mobility is high. In such scenarios it is possible that a node moves out of transmission range before sending the CTF (which nevertheless may suppress the CTFs of other nodes) or before receiving the actual data packet. Currently, the active selection scheme uses no recovery strategy that attempts to retransmit a packet if no CTF is heard after the timeout interval T , and the packet is lost.

In contrast to the CBF schemes and to the optimized greedy approach, the basic greedy




Fig. 12. Packet delivery ratio for scenarios with 300 nodes.

scheme performs significantly worse under mobility. With a maximum node speed of 50 m/s the packet delivery ratio drops to 0.2 with a beacon interval of 2 s. ‘Basic greedy’ selects a greedy forwarder out of the list of neighbors and tries to transmit the packet to it. If a neighbor moves out of transmission range, its entry expires and it is removed from the neighbor table after a timeout period during which no packets are received.¹⁰ During this period, all packets handed down to the link layer with this node as next hop are lost. The optimized greedy scheme detects these failures and reroutes all packets in the MAC queue destined for this next hop. Consequently, no packets are lost when the best suitable neighbor leaves the radio range if there is another suitable next hop in the neighbor table. The higher the node mobility, the more packets cannot be delivered with the basic greedy scheme and are therefore re-queued by the optimized scheme. Hence, the good performance of the optimized scheme comes at the expense of a trial-and-error strategy to detect a suitable forwarder that is still in transmission range, which may significantly increase the per hop delay (see also Section 5.5) and the network load. The CBF schemes achieve similar packet delivery ratios without any link layer packet loss recovery for the packet transmissions.

The same scenarios have also been simulated for densities of 100, 200, and 400 nodes within the 4 km² simulation area (not shown here). Generally, low node densities with only 100 or 200 nodes reduce the likelihood of greedy routes to the destination and all schemes achieve lower packet delivery ratios. With 400 nodes, the optimized greedy scheme, the basic CBF scheme, and the area-based CBF scheme deliver 100% of the packets. Active selection achieves a delivery ratio slightly below 100% with high mobility scenarios for reasons explained above. The performance of the the basic greedy schemes improves only marginally.

¹⁰ This beacon expiry timeout is usually a multiple of the beacon interval. We chose it as 3.5 times the beacon interval as in the simulations in [3].

5.4. Transmission costs

In Fig. 13 we show the transmission costs for the optimized greedy schemes and the CBF mechanisms in terms of average number of bytes transmitted at the MAC layer over the course of the simulation. The basic greedy schemes were omitted for lack of comparability; at high mobility, the packet delivery ratio is too low to allow a meaningful interpretation of the total overhead.

As expected, all CBF methods use less bandwidth than the greedy schemes together with the overhead caused by the beacon messages.¹¹ The area-based scheme consumes the least bandwidth, as no packet duplication occurs and—given a node density of approximately 15 nodes within transmission range—the forwarding node is almost always located within the Reuleaux triangle. Active selection causes a slightly higher overhead through the additional RTF and CTF messages and the basic CBF schemes causes the highest transmission costs due to packet duplication. The bandwidth consumption of all CBF schemes is relatively independent of mobility. The slight decrease in overhead can mainly be attributed to the decrease in the average path length caused by the random waypoint model.

The overhead caused by optimized greedy routing depends on a number of factors. The amount of data transmitted for beacon messages scales proportionally to the number of nodes, the beacon interval, and the simulation time. The value decreases somewhat with an increase in traffic since implicit beaconing causes beacons to be piggybacked on the data packets. Furthermore, the transmission costs for the greedy scheme increases significantly with an increase in mobility. The better the available neighbor information due to a high beacon rate, the lower the increase in MAC overhead caused by increasing mobility. When mobility is high, a large fraction of the packets have to be sent multiple times because of the MAC callback. This ratio decreases when more

¹¹ Results are significantly worse for the greedy schemes when we investigate the number of packets instead of the amount of bytes, since beacon messages are generally much smaller than data packets.




Fig. 13. Transmission costs on MAC layer for 300 nodes.




Fig. 14. Cost composition of greedy opt 2.0.

accurate neighbor information is available, at the expense of an increase of the overhead caused by the beacons.

To analyze the transmission costs caused by the optimized greedy scheme in more detail, Fig. 14 shows the specific components of MAC traffic for a beacon interval of 2 s and the scenarios with 300 nodes.

The bandwidth consumed by beacon messages and MAC control packets (i.e., unicast acknowledgments of the data packets) is independent of the mobility rate. In contrast, the overhead caused by the transmission of data packets increases significantly with higher mobility. Without mobility, optimized greedy consumes about as much bandwidth as area-based CBF. For a maximum node speed of 30 m/s, optimized greedy already consumes the same bandwidth as the active selection scheme (while the additional RTF/CTF messages in the active selection scheme also provide protection against the hidden terminal problem). For node speeds of 50 m/s and above, the greedy scheme even significantly exceeds the bandwidth usage of the basic CBF scheme with its unsuppressed duplicates. At this node mobility, the forwarding overhead is higher than the overhead caused by the beacon messages of all 300 nodes and exceeds the forwarding overhead with no mobility almost by a factor of four.

With only one sender and receiver and a data rate of 4 KBit/s, the amount of data traffic is ex-

tremely low given the total number of nodes. At such low rates, the additional traffic caused by the optimized greedy scheme can be handled by the MAC layer without any problems. However, for reasonable combinations of beacon traffic and actual data traffic, we expect the overhead ratio to become much worse. When the additional traffic caused by repeated MAC callback results in congestion, data packets as well as beacon messages may be lost. The former have to be retransmitted at the cost of additional bandwidth consumption, while loss of the latter decrease the accuracy of the neighbor tables, further aggravating the MAC callback problem.

5.5. Forwarding delay

For all CBF simulations, the maximum response time T was set to 45 ms. This parameter has a large impact on the average latency and was not subject to optimization. The optimal setting of T depends to a large degree on the MAC protocol and can be significantly reduced by integrating MAC and CBF. The parameter should further be dynamically adjusted to the node density and to network load. An optimized maximum response time adjustment strategy is left for future work.

Nevertheless, an analysis of packet forwarding latencies confirms the observations regarding the protocol overhead. Fig. 15 shows the average per hop latency (i.e., the time required by a packet to




Fig. 15. Average hop latency for 300 nodes.

travel from source to destination divided by the average number of hops of the route). Comparing the CBF schemes, the basic scheme has the lowest latency. There is no RTF/CTF handshake as with active selection and no sequential querying of regions as in the area-based scheme. This also explains the delay characteristics of the other two CBF schemes: with zero mobility the select scheme performs slightly better, because sometimes better routes may be found than with the ordered querying of areas. In a static scenario, this affects all packets, causing a perceptible difference in latencies. Mobility alleviates this effect and area-based CBF achieves slightly lower latency values.

When comparing CBF with the optimized greedy strategy, one can observe a similar behavior as with the transmission overhead. While mobility even *reduces* the forwarding delay of the CBF schemes, the delay of optimized greedy increases drastically with higher mobility. For maximum mobility rates of more than 30 m/s, the forwarding delay is larger than that of the unoptimized CBF schemes. The responsibility for this effect lies again in the increasing number of link layer retransmissions.

6. Conclusions and outlook

The advantage of position-based routing over other ad hoc routing protocols is the fact that

nodes require only knowledge about the local neighborhood and the destination's location instead of global route topology. Therefore, position-based routing is better suited for networks with a certain degree of mobility. With the CBF mechanism proposed in this paper, even this local knowledge and hence the sending of beacon messages is no longer required. Any node with progress toward a destination can participate in the forwarding process without the need for this node to be registered in a neighbor table. For CBF, data packets are transmitted via single-hop broadcast. All nodes within radio range and with forward progress toward the destination are eligible to continue to forward the packet. Thus, the responsibility for the forwarding decision now lies with the set of possible next hops instead of the forwarding node, as is the case in conventional forwarding methods. Forwarding takes place after a contention period during which one or more nodes are selected as next hops. Selection of more than one next hop causes unwanted packet duplication. We presented different suppression strategies to avoid this.

For existing position-based forwarding schemes, node mobility results in frequent beacon messages to keep the neighbor tables reasonably up-to-date. Particularly for highly mobile networks, CBF can provide significant bandwidth savings through the elimination of beacon messages and the reduction of MAC layer retries for packet transmissions caused by inaccurate neighbor tables. Furthermore, the decrease in the total number of packets reduces the probability of packet collisions and inefficient routing caused by inaccurate neighbor tables is avoided.

The simulation results presented in this paper show that excessive re-sending of data due to outdated neighbor table entries as it is the case for traditional position-based routing can be completely avoided by the proposed CBF approach. Since CBF does not require any beaconing, and since CBF together with the area-based suppression strategy does not lead to any noticeable packet duplication, the resulting data volume overhead of the contention-based method is much less than the data volume overhead generated with traditional position-based routing in highly mobile

ad hoc networks. Clearly, reducing load on the wireless medium is beneficial for ad hoc networking in general. In the rare case where a packet duplication occurs due to CBF, a simple strategy exists to improve the proposed suppression schemes: if duplication of packets occurs, these packets will be routed to the same destination at roughly the same time. Even with a very small state about which packets were recently forwarded, the duplicates can easily be suppressed in later nodes.

Thus, packet duplication can be reduced while the simplicity of the suppression schemes is retained. In addition to the reduced forwarding overhead, the CBF schemes also provide a lower packet forwarding delay when node mobility is high. For the simulations, we used very conservative timer settings and we expect the reduction in forwarding delay to be much more pronounced with a well tuned CBF implementation.

One key item of future work will be the integration of CBF and MAC functionality. Since both serve a somewhat similar purpose their integration can significantly reduce the overhead incurred by the CBF scheme. In particular, we expect that it is possible to significantly reduce the runtime of the random timers used for the contention process. If a MAC layer with RTS/CTS is used to solve the hidden terminal problem (as is possible with IEEE 802.11), it can be combined with the RTF/CTF messages of active selection which will significantly increase the efficiency of this suppression strategy. Furthermore, a maximum response time T which adapts to network load and node density can reduce the delay incurred by the contention period. So far, we have only considered greedy forwarding. In position-based routing, greedy forwarding fails if no neighbor with progress toward the destination exists. In such a case, a recovery strategy is used to circumnavigate the area with no reception. While Geodesic Anchors, as proposed for Terminodes routing, is directly applicable to CBF we also plan to investigate other alternatives that are more similar to the repair strategies of face-2 and GPSR.

The use of directional antennas in ad hoc networks recently gained increased scientific interest [29]. This technology seems to be a promising

candidate particularly in the context of area-based suppression.

References

- [1] T.-C. Hou, V.O. Li, Transmission range control in multihop packet radio networks, *IEEE Trans. Commun.* 34 (1) (1986) 38–44.
- [2] H. Takagi, L. Kleinrock, Optimal transmission ranges for randomly distributed packet radio terminals, *IEEE Trans. Commun.* 32 (3) (1984) 246–257.
- [3] B.N. Karp, H.T. Kung, GPSR: greedy perimeter stateless routing for wireless networks, in: *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '00)*, Boston, MA, 2000, pp. 243–254.
- [4] S. Basagni, I. Chlamtac, V.R. Syrotiuk, B.A. Woodward, A distance routing effect algorithm for mobility (DREAM), in: *Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '98)*, Dallas, TX, 1998, pp. 76–84.
- [5] L. Blažević, S. Giordano, J.-Y. LeBoudec, Self-organizing wide-area routing, in: *Proceedings of SCI 2000/ISAS 2000*, Orlando, FL, 2000.
- [6] P. Bose, P. Morin, I. Stojmenovic, J. Urrutia, Routing with guaranteed delivery in ad hoc wireless networks, in: *Proceedings of the 3rd International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M '99)*, Seattle, WS, 1999, pp. 48–55.
- [7] R. Morris, J. Jannotti, F. Kaashoek, J. Li, D.S.J. DeCoito, CarNet: A scalable ad hoc wireless network system, in: *Proceedings of the 9th ACM SIGOPS European Workshop: Beyond the PC: New Challenges for the Operating System*, Kolding, Denmark, 2000, p. 127ff.
- [8] H. Hartenstein, B. Bochow, A. Ebner, M. Lott, M. Radimirsch, D. Vollmer, Position-aware ad hoc wireless networks for inter-vehicle communications: The FleetNet project, in: *Proceedings of the Second ACM International Symposium on Mobile and Ad hoc Networking & Computing (MobiHoc '01)*, Long Beach, CA, 2001.
- [9] D.B. Johnson, D.A. Maltz, Dynamic source routing in ad hoc wireless networks, in: T. Imielinski, H. Korth (Eds.), *Mobile Computing*, vol. 353, Kluwer Academic Publishers, Boston, MA, 1996, pp. 153–181.
- [10] C.E. Perkins, E.M. Royer, Ad hoc on-demand distance vector routing, in: *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, New Orleans, LA, 1999, pp. 90–100.
- [11] E.M. Royer, C.-K. Toh, A review of current routing protocols for ad hoc mobile wireless networks, *IEEE Personal Commun.* 6 (2) (1999) 46–55.
- [12] M. Mauve, J. Widmer, H. Hartenstein, A survey on position-based routing in mobile ad hoc networks, *IEEE Network* 15 (6) (2001) 30–39.

- [13] M. Heissenbüttel, T. Braun, A novel position-based and beacon-less routing algorithm for mobile ad hoc networks, in: Proceedings of the 3rd IEEE Workshop on Applications and Services in Wireless Networks (ASWN '03), Bern, Switzerland, 2003, pp. 197–209.
- [14] Y.-B. Ko, N.H. Vaidya, Location-aided routing (LAR) in mobile ad hoc networks, in: Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '98), Dallas, TX, 1998, pp. 66–75.
- [15] S. Giordano, M. Hamdi, Mobility management: the virtual home region, Tech. Rep. SSC/1999/037, EPFL-ICA, October 1999.
- [16] J. Li, J. Jannotti, D.S.J. DeCouto, D.R. Karger, R. Morris, A scalable location service for geographic ad hoc routing, in: Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '00), Boston, MA, 2000, pp. 120–130.
- [17] T. Camp, J. Boleng, L. Wilcox, Location Information Services in Mobile Ad Hoc Networks, in: Proceedings of the IEEE International Conference on Communications (ICC), New York, 2002, pp. 3318–3324.
- [18] J. Nonnenmacher, E.W. Biersack, Scalable feedback for large groups, *IEEE/ACM Trans. Networking* 7 (3) (1999) 375–386.
- [19] T.T. Fuhrmann, J. Widmer, On the scaling of feedback algorithms for very large multicast groups, *Computer Communications* 24 (5–6) (2001) 539–547.
- [20] ANSI/IEEE, ANSI/IEEE Std 802. 11, 1999th Edition, 1999.
- [21] L. Kleinrock, F.A. Tobagi, Packet switching in radio channels: Part I—carrier sense multiple-access modes and their throughput-delay characteristics, *IEEE Trans. Commun.* 23 (12) (1975) 1400–1416.
- [22] W. Gleißner, H. Zeitler, The Reuleaux triangle and its center of mass, *Results Math.* 37 (2000) 335–344.
- [23] P. Karn, MACA—a new channel access method for packet radio, in: Proceedings of the 9th ARRL/CRRRL Amateur Radio Computer Networking Conference, 1990, pp. 134–140.
- [24] M. Matsumoto, T. Nishimura, Mersenne Twister: a 623-dimensionally equidistributed uniform pseudo-random number generator, *ACM Trans. Model. Comput. Simul.* 8 (1) (1998) 3–30.
- [25] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, M. Booth, F. Rossi, GNU Scientific Library—Reference Manual, <http://sources.redhat.com/gsl>, January 2002.
- [26] The ns-2 network simulator, <http://www.isi.edu/nsnam/ns/>.
- [27] C. Bettstetter, Mobility modeling in wireless networks: categorization, smooth movement, and border effects, *ACM SIGMOBILE Mobile Comput. Commun. Rev. (MC2R)* 5 (3) (2001) 55–67.
- [28] M. Käsemann, Beaconless position-based routing for mobile ad hoc networks, Master's thesis, Department of Mathematics and Computer Science, University of Mannheim, 2003.

- [29] L. Bao, J.J. Garcia-Luna-Aceves, Transmission scheduling in ad hoc networks with directional antennas, in: Proceedings of the Eighth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '02), Atlanta, GA, 2002, pp. 48–58.



Holger Füllner received his M.S. degree in Business Administration and Computer Science (Diplom-Wirtschaftsinformatiker) from the University of Mannheim, Germany in 2001. Since then, he is a research staff member at the “Lehrstuhl für Praktische Informatik IV” of the University of Mannheim, working toward his Ph.D. degree in Computer Science. His work is funded by NEC Network Labs Europe within the framework of the FleetNet project by the BMBF (German Ministry of Education and Research). His research interest focusses on Mobile Ad Hoc Networks, especially between street-bound vehicles.



Jörg Widmer received his M.S. degree in Business Administration and Computer Science and his Ph.D. degree in Computer Science from University of Mannheim, Germany in 2000 and 2003, respectively. In 1999 and 2000, he was at ICIR, Berkeley, CA (formerly ACIRI) to write his M.S. thesis on “Equation Based Congestion Control”. Since 2003 he is employed as senior researcher at the Swiss Federal Institute of Technology—Lausanne (EPFL) where he is working on ultra-wide band wireless communication.

His research interests include ad hoc networking, congestion control, and multicast communication.



Michael Käsemann received a Master's degree in Technical Computer Science in 2003 at the University of Mannheim, Mannheim, Germany. He was a research intern at the NEC Network Laboratories Europe from 2002 to 2003, where he focused on Mobile Ad Hoc Networking and Network Simulations. His research interests include mobile communications, network routing, multimedia and hardware design.



Martin Mauve received his M.S. and Ph.D. degrees in Computer Science from the University of Mannheim, Germany, in 1997 and 2000, respectively. He was an assistant professor at the University of Mannheim from October 2000 until March 2003. Since April 2003 he is a professor at the Heinrich Heine University of Düsseldorf, leading the research group on Computer Networks and Communication Systems. His research interests include mobile ad hoc networks, inter-vehicle communications and distributed multimedia systems.



Hannes Hartenstein received a Master's degree in Mathematics in 1995 and the Ph.D. degree in Computer Science in 1998, both from Albert-Ludwigs-Universitaet, Freiburg, Germany. He was Erasmus Scholar with the University of East Anglia, Norwich, UK, in 1991–1992 and received the Capocelli Award from the IEEE Data Compression Conference 1997 (with M. Ruhl) for the paper 'Optimal fractal coding is NP-hard'. He joined NEC Network Laboratories Europe in 1999 and is now a Senior Research Staff

Member with the Mobile Internet Group. He currently focusses

on IP-based mobility management as well as on ad hoc routing. His research interests include mobile communications, networking, multimedia, security, and theoretical computer science.