

REIHE INFORMATIK
TR-03-004

Position-Based Multicast Routing for Mobile Ad-Hoc Networks

Martin Mauve Holger Füßler Jörg Widmer Thomas Lang

Universität Mannheim
– Fakultät für Mathematik und Informatik –
Praktische Informatik IV
L15, 16
D-68131 Mannheim, Germany

Position-Based Multicast Routing for Mobile Ad-Hoc Networks

Martin Mauve, Holger Füßler, Jörg Widmer, Thomas Lang
Praktische Informatik IV, University of Mannheim, Germany

{mauve,fuessler,widmer}@informatik.uni-mannheim.de, lang@rumms.uni-mannheim.de

Abstract—In this paper we present **Position-Based Multicast (PBM)**, a multicast routing algorithm for mobile ad-hoc networks which does neither require the maintenance of a distribution structure (e.g., a tree or a mesh) nor resorts to flooding of data packets. Instead a forwarding node uses information about the positions of the destinations and its own neighbors to determine the next hops that the packet should be forwarded to and is thus very well suited for highly dynamic networks. PBM is a generalization of existing position-based unicast routing protocols such as face-2 or GPSR. The key contributions of PBM are rules for the splitting of multicast packets and a repair strategy for situations where there exists no direct neighbor that makes progress toward one or more destinations. The characteristics of PBM are evaluated in detail by means of simulation.

Index Terms—Ad-Hoc Networks, Position-Based Routing, Multicast.

I. INTRODUCTION

Many applications envisioned for mobile ad-hoc networks rely on group communication. Communication during disaster relief, networked games, and emergency warnings in vehicular networks are common examples for these applications. As a consequence multicast routing in mobile ad-hoc networks has received significant attention over the recent years.

In this paper we present a Position-Based Multicast routing protocol (PBM), which uses the geographic position of the nodes to make forwarding decisions. In contrast to existing approaches PBM neither requires the maintenance of a distribution structure (i.e., a tree or a mesh) nor resorts to flooding. PBM is a generalization of existing position-based unicast routing protocols, such as face-2 [1] or Greedy Perimeter Stateless Routing (GPSR) [2]. The general idea of these position-based unicast routing algorithms is to select the next hop based on position information such that a packet is forwarded in the geographical direction of the destination.

Position-based routing can be divided into two main functional elements: the location service and position-based forwarding. The *location service* is used to map

the unique identifier (such as an IP address) of a node to its geographical position. 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' destinations. Examples for existing location services that can be used for this purpose are Homezone [3], the Grid Location Service (GLS) [4] or the location service part of DREAM [5].

Position-based forwarding for unicast 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, for the forwarding decision the geographical positions of the node itself, its direct neighbors, and the packet's destination need to be known. With this information, the forwarding node selects one of its neighbors as a next hop such that the packet makes progress towards the geographical position of the destination. It is possible that there is no neighbor with progress towards the destination while there still exists a valid route to the destination. 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 towards 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 for environments where the nodes have access to their geographical position, such as in inter-vehicle-communication [6], [7].

In order to extend position-based routing to multicast two key problems have to be solved. First, at certain nodes a multicast packet has to be split into multiple copies in order to reach all destinations, the challenge being to decide when such a copy should be created. Second, the recovery strategy used to escape from a local optimum needs to be adapted to take multiple destinations into account. The key contributions of this work are solutions for both prob-

lems. The proposed algorithms are evaluated by means of simulation.

The remainder of this work is structured as follows. Section II gives an overview of related work in the area of multicast routing for mobile ad-hoc networks. Section III describes the Position-Based Multicast protocol in detail. The characteristics and performance of the protocol are then investigated by means of simulations in Section IV. Section V outlines how membership and position information could be made available in a scalable manner while Section VI concludes this paper with a summary and an outlook to future work.

II. RELATED WORK

Most existing multicast routing protocols for mobile ad-hoc networks maintain some form of distribution structure for the delivery of multicast packets. They can be classified into tree-based and mesh-based approaches. In tree-based multicast routing protocols data packets are forwarded on a single path to a given receiver. The union of the paths to all receivers forms the multicast-tree, which may be sender specific or common to all senders in a multicast session. Examples of tree-based multicast routing protocols for mobile ad-hoc networks are: Reservation-Based Multicast (RMB) [8], Ad Hoc Multicast Routing Protocol (AMRoute) [9], Ad Hoc Multicast Routing Protocol Utilizing Increasing Id-Numbers (AMRIS) [10], Bandwidth Efficient Multicast Routing Protocol [11], Multicast Ad Hoc on Demand Distance Vector Routing (MAODV) [12], Multicast Core Extraction Distributed Ad Hoc Routing (MCEDAR) [13], and Adaptive Demand-Driven Multicast Routing (ADMR) [14]. Typically these approaches include mechanisms such as local repair of the distribution tree or backup paths in order to compensate for the frequent topological changes in an ad-hoc network.

In mesh-based approaches there may be multiple paths to each receiver. This redundancy provides increased protection against topological changes. Examples of mesh-based multicast routing protocols for mobile ad-hoc networks are: On Demand Multicast Routing Protocol (ODMRP) [15], Core-Assisted Mesh Protocol (CAMP) [16], Neighbor Supporting Ad hoc Multicast Routing Protocol (NSMRP) [17], and Dynamic Core Based Multicast Routing Protocol (DCMP) [18].

Both, tree- and mesh-based multicast routing protocols, need to maintain state information about the distribution structure. Thus they are limited to environments where the node mobility and the distance covered by the network is such that the state is still valid (or can be locally repaired) while a packet traverses the distribution structure.

When this cannot be guaranteed, it has been proposed to use flooding to achieve a reliable form of multicast [19].

Knowledge about the geographical position of nodes has been used in [20] to improve ODMRP with mobility prediction and in [21] to limit flooding when the multicast group members reside in one specific geographic area. In Dynamic Source Multicast (DSM) [22] each node floods the network with information about its own position, thus each node knows the position of all other nodes in the ad-hoc network. The sender of a multicast packet then constructs a multicast tree from the position information of all receivers. This tree is efficiently encoded in the header of the packet.

In contrast to the existing approaches Position-Based Multicast (PBM), as described in this work, does neither require the maintenance of state about a distribution structure nor does it resort to flooding of the data packets. Instead each node that forwards a multicast packet autonomously determines the neighbors that it should forward the packet to. This decision is based on information about the position of the destination nodes, the position of the forwarding node, and the position of the forwarding node's neighbors. It can be regarded as an adaptation of position-based unicast routing schemes such as face-2 [1] and Greedy Perimeter Stateless Routing (GPSR) [2] to multicast routing.

III. POSITION-BASED MULTICAST

For multicast it is necessary to establish a distribution tree among the nodes, along which packets are forwarded towards the destinations. At the branching points of the tree, copies of the packet are sent along all the branches. Two – potentially conflicting – properties are desirable for such a distribution tree: (1) the length of the paths to the individual destinations should be minimal and (2) the total number of hops needed to forward the packet to all destinations should be as small as possible. If the topology of the network is known, a distribution tree that optimizes the first criterion can be obtained by combining the shortest paths to the destinations. Wherever these paths diverge, the packet is split. The second criterion is optimized by so-called Steiner trees (see e.g., [23]) which connect source and destinations with the minimum possible number of hops. A formulation of the Steiner problem for wireless networks where packets are broadcast to neighboring nodes is given in [24]. However, with position-based routing, routing decisions are based solely on local knowledge, thus neither the shortest paths to all destinations nor (heuristics for) Steiner trees can be used directly. Instead PBM uses locally available information to approximate the optima for both properties.

For the remainder of this work we assume that each node that forwards a packet has access to the following information:

- 1) *The node's own geographical position*: This information can be provided by a positioning service such as GPS [25].
- 2) *The position of all neighbors within transmission range*: The position of a node is made available to its direct neighbors in form of periodically transmitted beacons.
- 3) *The positions of the destinations*: these may be included in the packet or available locally (i.e., because a location service distributes position information about all nodes to all other nodes within the network, such as it is done in DREAM).

Given this information the main task of a forwarding node in PBM is to find a set of neighbors that should forward the packet next. We call these neighbors the *next hop nodes*. The current node will assign each destination of the packet to exactly one next hop node. Each next hop node then becomes forwarding node for this packet towards the assigned destinations. If the current node selects more than one next hop node, then the multicast packet is split. This may be required in order to reach destinations which are located in different directions relative to the forwarding node. The most important property of PBM is that each forwarding node autonomously decides how to forward the packet. This decision requires no global distribution structure such as a tree or a mesh.

There are two distinct cases that can occur when a forwarding node selects the next hop nodes: either for each destination exists at least one neighbor which is closer to that destination than the forwarding node itself. In this case *greedy multicast forwarding* is used. Otherwise the node employs *perimeter multicast forwarding*.

A. Greedy Multicast Forwarding

As discussed above, a multicast distribution tree ideally optimizes two criteria. First, the distance towards the destination nodes should be minimized and hence the progress of the packet towards the destinations maximized. Second, the (global) bandwidth usage should be minimized. Thus the objective function of a forwarding node should consist of two elements, one for each objective. Optimizing the progress of the packet can be done in the following way. Let k be the forwarding node, N the set of all neighbors of k , W the set of all subsets of N , Z the set of all destination nodes, and $d(x,y)$ a function which measures the distance between nodes x and y . Given a set of next hop nodes $w \in W$ the overall remaining distance to all destinations of a multicast packet can be calculated as

shown in Equation 1. In this equation for each destination the next hop node in the set w is chosen which is closest to that destination. Using Equation 1 as the sole optimization criterion would lead to a splitting of the multicast packet as soon as there is no single neighbor which provides the largest progress towards all destinations. This may be undesirable since it ignores the bandwidth usage.

$$f_d(w) = \sum_{z \in Z} \min_{m \in w} (d(m,z)) \quad (1)$$

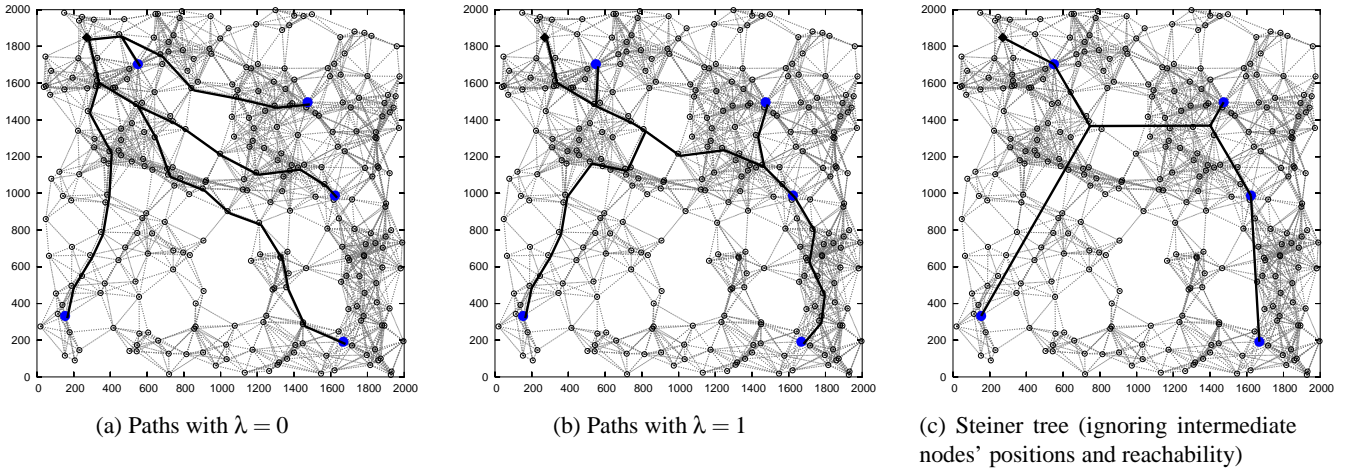
In order to consider the bandwidth usage we include the number of next hop nodes as a second element into the optimization criterion. The overall optimization criterion that determines which set of next hop nodes $w \in W$ should be selected as next forwarding nodes is given in Equation 2.

$$f(w) = \lambda \frac{|w|}{|N|} + (1 - \lambda) \frac{\sum_{z \in Z} \min_{m \in w} (d(m,z))}{\sum_{z \in Z} (d(k,z))} \quad (2)$$

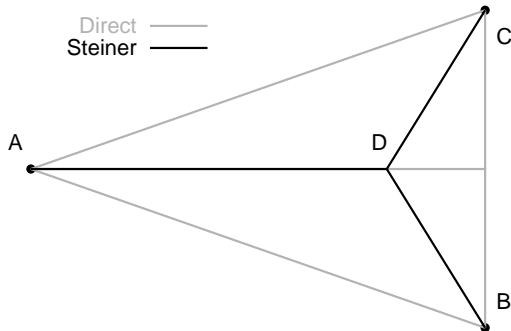
The first part of the equation determines the number of next hop neighbors and normalizes it to a value between $[0, 1]$ by dividing it by the total number of neighbors of k . The second part determines the remaining overall distance from the next hop nodes towards the destinations and normalizes this to a value between $[0, 1]$ by dividing it by the remaining overall distance calculated from the forwarding node k to the destinations. $\lambda \in [0, 1]$ determines the weight of each objective. If λ is close to 0 multicast packets will be split early, while for λ close to 1 the multicast packet will only be split if this is enforced by the restriction that there must be progress for each destination. An example for the impact of λ on the path that a multicast packet takes through the network is shown in Figure 1.

It is to be expected that the number of hops that a packet traverses from the source to a given destination increases with increasing λ , i.e., the path towards each destination becomes less direct. On the other hand the total number of single hop transmissions required to deliver the packet from the source to all destinations is likely to decrease when λ increases from 0 up to a certain value $s < 1$. The decrease of single hop transmissions when λ is increased from a value close to 0 to s is caused by the fact that packets are split later and thus less single hop transmissions take place. However, if a packet is split very late, i.e., $\lambda > s$ then the total number of hops may increase again.

These considerations can be illustrated with the simple topology given in Figure 2. Let A be the forwarding node and B as well as C the destinations of a multicast packet. Let us further assume that the node density is high enough so that a packet can be split virtually anywhere and that the

Fig. 1. Effect of λ

distance to the destinations is much larger than the radio range. If A decided to split the packet, the copies would be forwarded along \overline{AB} and \overline{AC} , taking the most direct path to the destinations. For a minimum number of total hops for the distribution of the packet to all destinations, the packet should be forwarded along \overline{AD} and node D should then split the packet and send copies to the final destinations (as indicated by “Steiner” in the graph). If the packet is not split at D but forwarded further, the total number of hops as well as the lengths of the individual paths to the destinations increase again. Therefore, the packet should ideally be split somewhere between A and D . Since λ determines how early a packet should be split, there will be a value $s < 1$ for λ where the total number of single hop transmission will be minimal. We determine a value for s by means of simulation in Section IV.

Fig. 2. Effect of λ on the number of single hop transmissions

B. Perimeter Multicast Forwarding

Applying greedy multicast forwarding may lead to a situation where the packet arrives at a node that does not have neighbors providing progress for one or more destinations. An example of this is depicted in Figure 3: the

copy of the multicast packet which is on its way to D_2 , D_3 , and D_4 , as well as the copy for D_5 get stuck in a local optimum.




Fig. 3. Greedy Multicast Routing Failure

For position-based unicast, this problem has been solved by applying a modification of the right hand rule ([1], [2]). The basic idea is to traverse the boundaries of gaps in the network until greedy forwarding can be resumed. To this end the graph formed by the connections (edges) between mobile nodes is planarized, i.e., intersecting edges are removed. This planarization is based on Relative Neighborhood or Gabriel Graphs [26], [27]. It can be done individually by each node based on local knowledge and does not partition the graph.

On the planarized graph the right hand rule can be used to escape the local optimum: the node where the local optimum is reached calculates a virtual edge from itself to the destination. The packet is then transmitted over the next edge counter-clockwise of that virtual edge. A packet transmitted this way is said to be in *perimeter mode*. When a packet is received by a node in perimeter mode, then this node checks if it is closer to the destination than the node where the packet entered perimeter mode. If this is the case the packet is reverted to *greedy mode* and forwarded in greedy fashion. If this is not the case the packet is forwarded over the next edge counter-clockwise from the edge it arrived on. The combination of perimeter and greedy forwarding guarantees that the destination is reached, as long as the network is static and as long as a valid connection between source and destination exists.

For PBM we generalized this algorithm to support packets with multiple destinations. If a node in PBM detects that it has no neighbors with forward progress for one or more destinations, then multicast perimeter mode is initialized for these destinations. For all other destinations greedy multicast forwarding is used. As in the unicast case the parameter mode is performed on the planarized graph (PBM uses Gabriel Graphs for planarization). The virtual edge used for the initialization is calculated as the connection between the current node and the position representing the average of the positions of the affected destination nodes. The multicast perimeter packet is then transmitted over the first edge counter-clockwise of the virtual edge.

When a node receives a perimeter multicast packet, it checks for each destination, if it is closer to that destination than the node where the packet entered perimeter multicast mode. For all destinations where this is the case greedy multicast forwarding can be resumed, for all other destinations perimeter multicasting is continued by transmitting the packet over the next edge counter-clockwise of the edge where the packet arrived.

Automatically splitting a packet into copies that are to be forwarded in greedy multicast mode and a copy that is to use perimeter multicast may cause the transmission of the same packet to two nodes which are located in the same direction, or even to the same node twice. In order to reduce the load on the network PBM includes an optional combination of greedy and perimeter multicast forwarding: if some, but not all, destinations of a packet require perimeter multicast forwarding, then the next hop is determined using the perimeter rules from above. All copies of the packet with destinations for which greedy forwarding could be used also select this node as the next hop, if

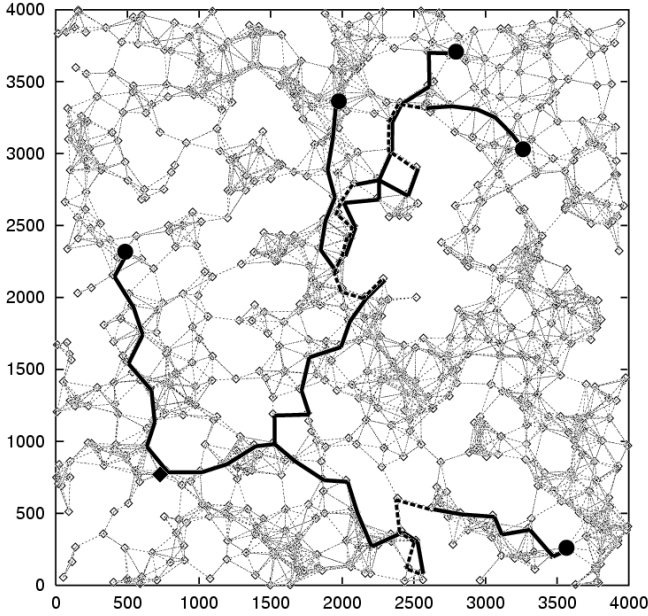
it provides progress towards the copy's destination. This reduces the number of copies of the same packet in the network. It comes at the cost of a potentially increased path length towards the individual destinations. Figure 4 shows how the problem depicted in Figure 3 is solved using perimeter multicast routing with and without combining perimeter and greedy packets.

IV. EVALUATION


We evaluated the performance and behavior of PBM by means of simulation. The primary goal of these simulations was to understand the proposed routing algorithm with acceptable empirical significance when used in network topologies of reasonable size. The simulation of multicast routing protocols is more demanding than the simulation of unicast routing for two main reasons:

- The simulated network needs to be sufficiently large to be able to distinguish between flooding and multicasting a packet. Given a radio range of 250 meters and 5 or more destinations, areas of the size commonly used for unicast simulations (with a side length of a few hundred meters to one or two kilometers) are too small for this purpose.
- The early loss of a multicast packet will lead to a drop for multiple destinations. Thus the loss rate is subject to a much higher variance than for unicast. As a consequence the number of simulation runs should be higher.

For these reasons the simulated networks contain more than 1000 nodes on an area of 4km by 4km or larger and the number of simulation runs is on the order of 1000 for a large number of parameter combinations (about 200). This is currently feasible neither with ns-2 nor with GloMoSim. We thus decided to implement the simulation in C++ without the use of a dedicated simulation environment. In our simulation nodes can communicate if they are within radio range, the transmission of a packet takes 10 ms and there is no simulation of a MAC-layer. We are fully aware that this does not allow the investigation of MAC-layer interaction with PBM. However, this approach had the key advantage that we were able to observe the characteristics of the routing algorithm itself in a much more detailed manner and with much more empirical significance than could otherwise be done. We intend to perform separate studies on the interaction of PBM with specific MAC layers such as IEEE 802.11 for selected parameter combinations and reduced network size in the future.



(a) Paths taken without combining perimeter and greedy packets



(b) Paths taken with combining perimeter and greedy packets

Fig. 4. Perimeter multicast routing

A. Simulation Setup

We simulated the behavior of PBM using three different simulation areas: small (2000 by 2000 meters), medium (4000 by 4000 meters), and large (8000 by 8000 meters). For each area we investigated multiple node densities (30, 40, 50, 60 node per km^2) with the nodes initially being randomly placed in the simulation area with an equal distribution. Node movement follows the random waypoint model [28], the node speed was randomly chosen with an equal distribution for each node out of an interval between 0m/s and a certain maximum speed (0m/s, 10m/s, 20m/s, 30m/s, 40m/s, 50m/s), the pause time was set to 0 seconds. For each packet one sender and a number of receivers (5, 10, 15, 20, 25, 30) was chosen such that all destinations reside in the same network partition as the sender. Then one packet was transmitted. After the packet traversed the network, the nodes were redistributed, and a new sender as well as new receivers were selected. This process was repeated 1000 times.

B. Delivery Rate

As in position-based unicast, PBM is guaranteed to successfully deliver all packets in a static network where the sender and all receivers reside within the same network partition. In a dynamic network the use of the perimeter mode may lead to routing loops and thus to packet drops. Figure 5 shows how such a loop can come into existence. In this figure the source of the packet is S , the destination is D and in node u the packet enters perimeter mode.

While the packet traverses the link from v to w , a connection is established between x and v because of node movement. The packet will then be caught in the triangle formed by v , w and x and will consequently be dropped.




Fig. 5. Routing Loop in a Dynamic Network

We investigated the likeliness of packet loss caused by this event with respect to mobility and node density. Only those simulation runs were taken into account where the sender and all receivers resided within the same partition for the complete simulation run. We counted the number of destinations that were not reached and related it to the overall number of destinations. The result is the loss rate. Figure 6 shows the loss rate for the medium size area with 5 destinations per transmitted packet. It can be seen that the likeliness for a packet drop caused by a routing loop increases with a decrease in node density. This is the case

since routing loops can only occur in perimeter mode and the likeliness for a packet using the perimeter mode increases with a decrease in node density. Also it can be observed that the likeliness for a routing loop increases when the node mobility increases. This is not surprising, since node mobility is the reason why a routing loop is formed. Examining the values of the loss rate, it can be noted that it remains fairly low (below 2%) for node densities above 50 nodes per km^2 , even if the node mobility is extremely high.




Fig. 6. Lossrate without Hop Limit

The second problem that perimeter routing may encounter is caused by the fact that the border of an area without nodes is always traversed counter-clockwise, even though half of the time a clockwise traversal would lead to a shorter route. If the area traversed in this way is very large, or if it is the outer boundary of the network, then the required number of hops for this traversal may be unacceptably high if the wrong choice for the orientation of the traversal is made. An example for this is depicted in Figure 7.

In order to determine the effect of this problem we assigned each packet at hop-count. When this hop-count exceeds a predefined value, then the packet is dropped. This value was set to 200 which prevents a packet from traversing the outer boundary of the network. Any packet exceeding this hop count was dropped. This was done in addition to the packet drops reported above. The result of this simulation is shown in Figure 8. It should be pointed out that this figure does include the drops caused by looping packets. With this fact it is remarkable that the total amount of lost packets is almost completely independent of node speeds. It is easy to see that the likeliness of encountering a perimeter that leads to a traversal of the boundaries of the network depends only on the node density. However, at a first glance one would expect that this is in addition to the packet drops caused by routing




Fig. 7. Perimeter Problem in a Static Network

loops. This is not the case since packets that traverse the boundary of the network have a much longer path than other packets. Thus the likeliness that they encounter a routing loop is much higher than for other packets. As a consequence, the vast majority of packets that are caught in routing loops are packets that traverse the boundary of the network. An increase of speed does therefore only change the reason why a packet is discarded (routing loop vs. hop count exceeded) but has no significant impact on the overall loss rate in Figure 8.




Fig. 8. Lossrate with a Hop Limit of 200

Concluding it can be said that for node densities of about 50 nodes or more per km^2 PBM will have very low loss rates. In addition the main cause for packet loss is the traversal of the network boundary. If the network boundary is sufficiently far away from communication partners then the loss rates should decrease substantially. Furthermore it seems worthwhile to investigate approaches to

improve the decision about the orientation of the traversal if information is available about the boundaries of the network (e.g., rivers, lakes). Since PBM is a generalization these observations also hold for position based unicast routing as proposed in face-2 and GPSR.

C. Average Path Length

We define the *average path length* to be the average number of hops that a packet traverses on its path from the sender to each receiver. Thus the average path length measures how direct the path towards the destinations is and thereby how much delay the packet will encounter. We were interested in understanding how the choice of λ would influence the average path length. Our hypothesis was that the average path length increases with increasing λ , since a small value for λ would lead to an optimization of the packet progress, while a large value of λ would delay the splitting of the packet. We varied λ from 0 to 1 in 0.05 increments for all combinations of the remaining simulation parameters as described above. We considered only those simulation runs which did not include packet loss or packets which traversed to outer boundary of the network.

Figure 9 shows how the path length depends on λ for the medium size region. This figure contains 24 graphs, each representing one combination of parameters: 40, 50 and 60 nodes per km^2 , maximum speed of 0, 20, 40, 60 and the optional combination of greedy and perimeter packets turned on and off. There are three groups of graphs, one group for each node density. It is clear that with an increase in node density the path length will decrease, since a more direct path becomes possible. All 24 graphs show the same main characteristic: the path length increases steadily while the value for λ is increased.

Surprisingly the combination of greedy and perimeter multicast packets did not have a major impact in any of our simulation runs. A further investigation suggested that it rarely alters the path of the packet significantly. As one would expect, the maximum speed had no impact on the path length.

D. Number of Single-Hop Transmissions

The *number of single-hop transmissions* is determined by counting all transmissions that are required to forward the multicast packet to all destinations. It is a measure for the load on the network, caused by the multicast packet. As described in the previous section we expected that for a given set of parameters the number of single-hop transmission will reach a global minimum for a value of λ between 0 and 1.




Fig. 9. Effect of λ on the average path length

Figure 10 shows how the number of single-hop transmissions depends on λ for the medium size region. Again the node density has a major impact: the more nodes, the less single-hops transmissions are required. This results in the same grouping of graphs as it has been observed for the path length. Over all simulations the shape of the graph is almost identical for all parameter combinations, with the minimum between 0.3 and 0.6. Neither the maximum speed nor the combination of greedy and perimeter packets did have a significant impact on the number of single-hop transmissions.

These results indicate that a true trade off between the goals of minimizing the average path length and minimizing the number of single-hop transmissions only exists for values of λ between 0 and 0.6. Values that are not in this range are dominated by the other values and should not be used.

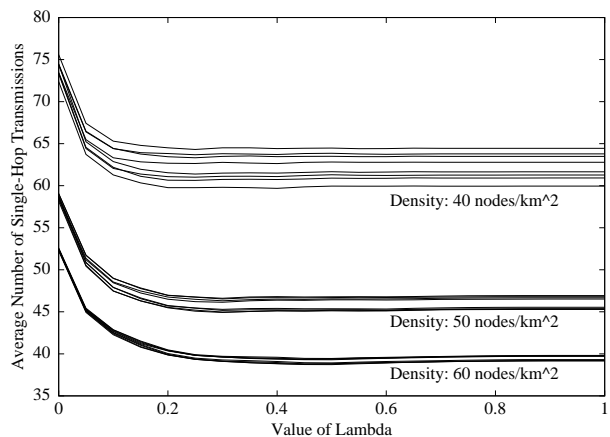


Fig. 10. Effect of λ on the number of single-hop transmissions

E. Bandwidth Reduction

Multicast is primarily used to reduce the bandwidth requirement when the same packet needs to be delivered to

multiple destinations. Thus it is interesting to compare the number of single hop transmissions that are required for the transmission of the packet when unicast is used to the number of single hop transmissions that are used when the packet is delivered via multicast. In order to make this comparison possible we determined the average path length with λ set to 0. We then multiplied this value by the number of destinations. This results in the number of single hop transmissions that would be required if position-based unicast had been used. We compared this value to the number of single hop transmissions of PBM with λ selected such that the number of single hop transmissions is minimal. This was done by dividing the multicast value by the unicast value for distinct settings of area size, number of nodes and number of destinations. Figure 11 shows how the reduction in single hop transmissions increases as the number of destinations grows. The setting from which this graph was derived is a medium sized area with a node density of 60 nodes per km^2 and a maximum node velocity of 30 meters per second. Other combinations of parameters yield similar results with the reduction of single-hop transmissions reaching about 66% for 30 destinations.

In addition to reducing the number of single hop transmission the usage of multicast also prevents the overload of the network close to the sender. These hot-spots appear at the sender if the same unicast packet is transmitted once per destination. While the overall load on the network is reduced by the amount shown in Figure 11 the reduction in those critical areas of the network is actually much higher: it is reduced by a factor which depends linearly on the number of destinations.

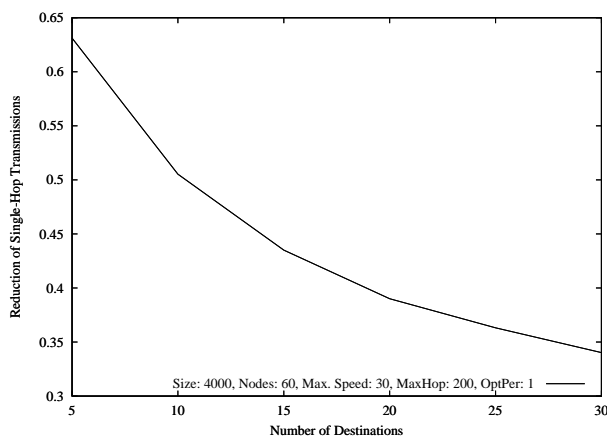


Fig. 11. Reduction of Single-Hop Transmissions

V. GROUP MEMBERSHIP AND POSITION INFORMATION

In order to work as described in the previous sections PBM requires that a forwarding node knows the identity

and the position of all destinations. While it is conceivable that the sender of a packet gathers this information and places it in the packet header, this does not seem to be viable for large receiver sets. In particular this would increase the size of the header and thus limit the key benefit of multicast, i.e., the reduction of the required bandwidth. We are currently investigating the modification of existing location services such as GLS [4] or DREAM [5] to include group membership information along with position information. This would allow PBM to make the forwarding decision at each node without including overhead in the data packets. Furthermore it is possible to aggregate multiple destinations that are located in one geographic region, such that the distribution of location and membership information requires only minimal resources. This aggregation can be done hierarchically such that more detailed information about the membership and position of members becomes available as the packet moves closer to those members.

VI. CONCLUSIONS

In this paper we presented a multicast routing algorithm for mobile ad-hoc networks. Position-Based Multicast (PBM) is a generalization of existing unicast routing algorithms (e.g., face-2 or GPSR) which use the geographic position of the participating nodes for the forwarding of packets. PBM consists of a greedy forwarding part that selects the next hop(s) of a packet based on the positions of the forwarding node, its neighbors and the destinations. Furthermore a recovery strategy is specified for situations where greedy forwarding fails. The key advantage of PBM is that no distribution structure like a tree or mesh needs to be constructed and maintained. Thus PBM is very well suited for highly dynamic networks without resorting to flooding of the data packets. In addition the rule for splitting a multicast packet includes a parameter λ that may be used to adapt the algorithm to different application scenarios by controlling the tradeoff between latency and bandwidth. The application of PBM to a large number of different network parameters (node speed, network size, node density) has been investigated by means of simulation. As a consequence the theoretical behavior of PBM in terms of drop rate, potential for bandwidth reduction and the effect of the parameter λ is very well understood. The simulation of selected scenarios with realistic MAC protocols is currently under way. The key issue that remains open is the scalable distribution of group membership and position information. We are currently working on a solution for this problem which is based on an integration of group membership information into existing location services.

REFERENCES

- [1] Prosenjit Bose, Pat Morin, Ivan Stojmenovic, and Jorge 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, August 1999, pp. 48–55.
- [2] Brad Karp and 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, Massachusetts, August 2000, pp. 243–254.
- [3] Silvia Giordano and Maher Hamdi, "Mobility Management: The Virtual Home Region," Tech. Rep. SSC/1999/037, EPFL-ICA, October 1999.
- [4] Jinyang Li, John Jannotti, Douglas S. J. DeCouto, David R. Karger, and Robert 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, Massachusetts, August 2000, pp. 120–130.
- [5] Stefano Basagni, Imrich Chlamtac, Violet R. Syrotiuk, and Barry 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, Texas, October 1998, pp. 76–84.
- [6] Robert Morris, John Jannotti, Frans Kaashoek, Jinyang Li, and Douglas S. J. DeCouto, "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, September 2000, p. 127ff.
- [7] "The FleetNet project," <http://www.fleetnet.de>.
- [8] M. S. Corson and S. G. Batsell, "A reservation-based multicast (RBM) routing protocol for mobile networks: Overview of initial route construction," in *Proc. of INFOCOM*, 1995, pp. 1063–1074.
- [9] E. Bommaiah, M. Liu, A. McAuley, and R. Talpade, "AM-Route: Ad hoc Multicast Routing Protocol," Internet-Draft, draft-talpade-manet-amroute-00.txt, work in progress, August 1998.
- [10] C. Wu, Y. Tay, and C. Toh, "Ad hoc Multicast Routing Protocol Utilizing Increasing id-numbers (AMRIS) Functional Specification," Internet-Draft, draft-ietf-manet-amris-spec-00.txt, work in progress, November 1998.
- [11] T. Ozaki, J. Kim, and T. Suda, "Bandwidth Efficient Multicast Routing Protocol for Ad hoc Networks," in *Proc of IEEE ICCCN*, October 1999, pp. 10–17.
- [12] E. Royer and C. Perkins, "Multicast using ad hoc on demand distance vector routing," in *Proc. of ACM MOBICOM*, 1999, pp. 207–218.
- [13] P. Sinha, S. Sivakumar, and V. Bharghavan, "MCEDAR: Multicast Core Extraction Distributed Ad hoc Routing," in *Proc. of IEEE WCNC*, August 1999, pp. 1313–1317.
- [14] J. Jetcheva and D. Johnson, "Adaptive demand-driven multicast routing in multi-hop wireless ad hoc networks," in *Proc. of MOBIHOC*, 2001, pp. 33–44.
- [15] S. Lee, M. Gerla, and C. Chiang, "On Demand Multicast Routing Protocol," in *Proc of IEEE WCNC*, August 1999, pp. 1298–1302.
- [16] J. Garcia-Luna-Aceves and E. Madruga, "The Core-Assisted Mesh Protocol," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 8, pp. 1380–1394, August 1999.
- [17] S. Lee and C. Kim, "Neighbor Supporting Ad hoc Multicast Routing Protocol," in *Proc. of ACM MOBIHOC*, August 2000, pp. 37–50.
- [18] S. Das, B. Manoj, and C. Murthy, "A Dynamic Core Based Multicast Routing Protocol for Ad hoc Wireless Networks," in *Proc. of ACM MOBIHOC*, June 2002.
- [19] K. Obraczka, K. Viswanath, and G. Tsudik, "Flooding for Reliable Multicast in Multi-Hop Ad Hoc Networks," *Wireless Networks*, vol. 7, no. 6, pp. 627–634, November 2001.
- [20] S. Su and M. Gerla, "Wireless ad hoc multicast routing with mobility prediction," *Mobile Networks and Applications*, vol. 6, no. 4, pp. 351–360, 2001.
- [21] Y. Ko and N. Vaidya, "Geocasting in Mobile Ad Hoc Networks: Location-Based Multicast Algorithms," in *Proc. of IEEE WMCSA*, February 1999.
- [22] I. Chlamtac S. Basagni and V. Syrotiuk, "Location aware, dependable multicast for mobile ad hoc networks," *Computer Networks*, vol. 36, no. 5-6, pp. 659–670, August 2001.
- [23] F. K. Hwang, D. S. Richards, and P. Winter, "The Steiner tree problem," *Annals of discrete mathematics*, vol. 53, 1992.
- [24] J. L. Wong, G. Veltri, and M. Potkonjak, "Energy-efficient data multicast in multi-hop wireless networks," in *Integrated Management of Power Aware Communications, Computing and Networking (IMPACCT)*, May 2002.
- [25] Elliot B. Kaplan, *Understanding GPS*, Artech House, 1996.
- [26] Godfried Toussaint, "The relative neighborhood graph of a finite planar set," *Pattern Recognition*, vol. 12, no. 4, pp. 261–268, 1980.
- [27] K. Gabriel and R. Sokal, "A new statistical approach to geographic variation analysis," *Systematic Zoology*, vol. 18, pp. 259–278, 1969.
- [28] Josh Broch, David A. Maltz, David B. Johnson, Yih-Chun Hu, and Jorjeta G. Jetcheva, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols," in *Proceedings of the fourth annual ACM/IEEE International Conference on Mobile computing and networking (MobiCom '98)*, Dallas, Texas, October 1998, pp. 85–97.