REIHE INFORMATIK
3/2002

A Comparison of Routing Strategies for Vehicular Ad-Hoc Networks

Holger Füßler, Martin Mauve
Universität Mannheim
Praktische Informatik IV
L15, 16
D-68131 Mannheim

Hannes Hartenstein, Michael Käsemann
NEC Network Laboratories
NL-E Heidelberg
Adenauerplatz 6
D-69115 Heidelberg

Dieter Vollmer
DaimlerChrysler AG
Research & Technology
Mercedesstraße 137
D-70546 Stuttgart

A Comparison of Routing Strategies in Vehicular Ad-Hoc Networks

Holger Füßler* Martin Mauve* Hannes Hartenstein† Dieter Vollmer‡
Michael Käsemann†

March 2002

Abstract

In this paper we investigate the use of ad-hoc routing algorithms for the exchange of data between vehicles. There are two main aspects that are of interest in this context: the specific characteristics of ad-hoc networks formed by vehicles and the applicability of existing ad-hoc routing schemes to networks that display these characteristics. In order to address both aspects we generate realistic vehicular movement patterns of highway traffic scenarios using a well validated traffic simulation tool. Based on these patterns we show that the characteristics of vehicular ad-hoc networks are quite different from the frequently used random waypoint model. We then proceed to evaluate the performance of a reactive ad-hoc routing protocol (DSR) and of a position-based approach (greedy forwarding as done in GPSR) in combination with a simple reactive location service. Our analysis suggests that for vehicular networks where communication spans more than 2 or 3 hops position-based ad-hoc routing has significant advantages over reactive non-position-based approaches both in the number of successfully delivered packets and in routing overhead.

1 Introduction

Communication between vehicles is considered a prime area where mobile ad-hoc-networks are likely to be deployed in the near future. The reasons for this are twofold. First, vehicles can easily provide the required power for wireless communication, and adding some weight for antennas and additional communication hardware does not cause major problems. Furthermore it can be expected that vehicles will have an accurate knowledge of their own geographical position, e.g., by means of GPS. Thus many problems making the deployment of ad-hoc networks in other scenarios problematic are not relevant here. Second, there is a wealth of desirable applications for ad-hoc communication between vehicles ranging from emergency warnings and distribution of traffic as well as road condition information to chatting and distributed games. As a consequence many vehicle manufacturers and their suppliers are actively supporting research on how to integrate mobile ad-hoc networks into vehicles [1, 2].

While communication between vehicles is frequently mentioned [3] as a target for ad-hoc routing protocols, there have previously been no studies on how the specific movement patterns of vehicles may influence the protocol performance and applicability. Typically the behavior of routing protocols for mobile ad-hoc networks is analyzed based on the assumption that the nodes in the network follow the random waypoint mobility model [4]. In this model each node randomly selects a waypoint in the area that contains the network and moves from its current location to the waypoint with a random but constant speed. Once a node has arrived at the waypoint it pauses for a random amount of time before selecting a new waypoint. Since this movement pattern of nodes has no similarity to the behavior of vehicles, the random waypoint model seems to be inappropriate to investigate the characteristics of vehicular ad-hoc networks or to determine which routing protocols are suitable for vehicular ad-hoc networks.

---

[*] Holger Füßler and Martin Mauve are with the "Lehrstuhl für Praktische Informatik IV" of the University of Mannheim ({fuessler, mauve}@informatik.uni-mannheim.de).

[†] Hannes Hartenstein and Michael Käsemann are with the NEC Network Laboratories Heidelberg ({Hannes.Hartenstein, Michael.Kaesemann}@ccrle.nec.de).

[‡] Dieter Vollmer is with DaimlerChrysler … (Dieter.Vollmer@daimlerchrysler.com).

In this paper we propose to make use of a detailed model for vehicular traffic to investigate the use of ad-hoc routing algorithms for the exchange of data between vehicles. The model includes elements such as vehicle characteristics (e.g., a car has a different movement pattern than a truck) and driver behavior (e.g., when does a driver decide to change lanes). Models like this are used by vehicle manufacturers to determine the lifetime of parts of a vehicle, such as shock absorbers or turn signals. Thus, these models have to be very accurate. As an output highly realistic movement patterns are produced.

Based on these movement patterns we are then able to analyze the characteristics of the dynamic topology formed by the mobile nodes. We show that the characteristics this network is significantly different from those formed by using the random waypoint model. In particular we look at network partitioning aspects and want to understand whether oncoming traffic needs to be used for the routing of packets. We employ the network simulator ns-2 [5] and the accompanying Ad-Hockey tool [6] for this purpose.

Our focus is then on studying the applicability of two routing strategies to vehicular ad-hoc networks by means of ns-2 and the above mentioned movement patterns. The key question we want to answer is whether the use of positional information in a routing approach provides significant benefits for this kind of network.

As a representative of a reactive non-position-based strategy we investigate the behavior of the well known Dynamic Source Routing (DSR) [7] protocol. As position-based strategy we build on the Greedy Perimeter Stateless Routing (GPSR) [8] protocol. Most of the previous studies on position-based routing are based on the assumption that an 'ideal' location service is present which has an accurate knowledge of the positions of all nodes at no cost. In order to prevent this unfairness from affecting our results, we implemented a simple reactive location service inspired by the DSR route discovery procedure.

The main contributions of this work are (1) a detailed discussion of the characteristics of vehicular ad-hoc networks (2) the comparison of ad-hoc routing strategies for vehicular networks by means of realistic movement patterns and (3) the use of a location service in combination with position-based routing such that the comparison is fair in the sense that neither approach profits from perfect external information.

The remainder of this work is structured as follows: we outline the model for vehicular movement in Section 2 where we also describe the generated movement pattern for a highway scenario used in the remainder of this work. The characteristics of vehicular ad-hoc networks are investigated in Section 3. In Section 4 we provide an overview of ad-hoc routing strategies and related work in general and give a more detailed description of the two candidate protocols we selected for the comparison. In Section 5 we report on the results of the simulation of both routing protocols for the communication between vehicles. We then conclude our findings and report on directions of future research in Section 6.

2 Simulation of Vehicular Traffic

Vehicular traffic simulations can be classified coarsely into *microscopic* and *macroscopic* approaches [9].

When following a macroscopic approach, one focuses on system parameters like *traffic density* (number of vehicles per kilometer per lane) or *traffic flow* (number of vehicles per hour crossing an intersection) in order to compute a road's capacity or the distribution of traffic in a road net. In general, from a macroscopic perspective vehicular traffic is viewed as a fluid compressible medium and, therefore, is modeled as a special derivation of the Navier-Stokes equations.

In contrast, with a microscopic approach the movement of *each* individual vehicle is determined. In order to generate vehicle movement patterns for ad hoc routing experiments one clearly has to follow a microscopic approach, since the position of each individual vehicle is needed. Nevertheless, one also has to take care that a microscopic simulation does not result in unrealistic macroscopic effects. As the vehicle movements are generated by a 'pre-process' and complexity is therefore a minor concern, we decided to use a *Driver Behavior Model* [10, 11] for the microscopic traffic simulation. Such a model not only takes the characteristics of the cars into account but it also includes a model of the driver's behavior, like lane changing and passing decisions, traffic regulation and traffic sign considerations, or decreasing speed in curves, to name only a few. Driver Behavior Models are known to be highly accurate and are therefore used by vehicle manufacturers, e.g., to determine the lifetime of certain parts of the car.

As a simulator we use the well validated DaimlerChrysler-internal driver behavior simulation tool called FARSI. This simulator is regularly employed to generate traffic simulations for the product

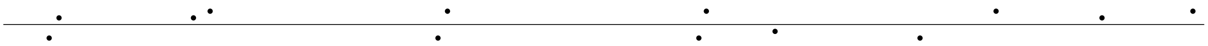

Figure 1: A 500m highway segment with a traffic density of 6 vehicles per kilometer and lane taken from our generated movement scenario.

development and evaluation of (company name). In particular FARSI simulations show realistic speeds, distances, and macroscopic properties like traffic flow and lane usage. Thus, FARSI guarantees that the vehicle movement patterns forming the basis of our experiments are as realistic as possible.

In this paper we investigate a typical highway scenario of 30 km length with two lanes per direction and with an average of 6 vehicles per kilometer and lane. Furthermore, the so-called 50%-desired speed parameter v_f (the parameter v_f splits the the population of vehicles into two halfs: the ones with a desired speed of at most v_f and the ones with a desired speed larger than v_f) is set to 130km/h. We assume that 15% of all vehicles are trucks. In FARSI the oncoming traffic is generated as a separate simulation for a single direction, i.e., both directions are independent. The positions of the vehicles are recorded every half a second together with current speed, lane identifier, and acceleration. From this file we generated our ns-2 movement file by taking a 200 seconds slice of the scenario.

The described scenario corresponds to weak day traffic on a German highway. In order to get an impression of the topology of a highway scenario with such a traffic density, a snapshot with realistic proportions for a highway segment of 500 m is given in Figure 1.

Since the topology and the topological changes over time are of utmost importance for our routing experiments, we present in the following some properties of the generated scenario with respect to the distribution of velocities and lane usage.

Figure 2 shows the distribution of the initial desired speeds for the simulation (we quantized speeds into 10km/h bins). The corresponding cumulative distribution shown in Figure 3 matches very well the cumulative distribution taken from a 'real' measurement from a German highway.

The distribution of velocities in the simulated scenario with 6 vehicle per km and lane is given in Figure 4 (percentage of total time spent in a specific velocity class).

The lane usage measurement for our highway scenario shows 57.2% usage of the left lane and 42.8% usage of the right lane. This is typical for weak day traffic on a German highway since vehicles are only allowed to pass on the left lane.[*]

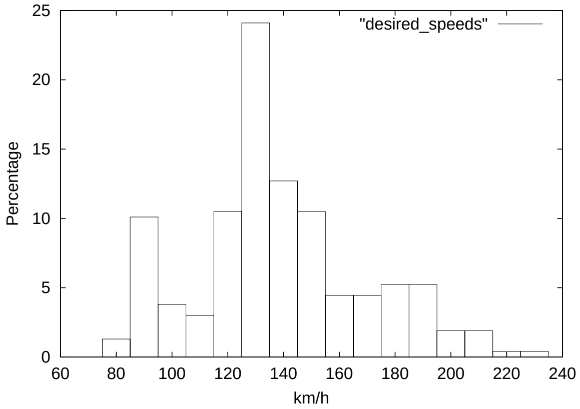

Figure 2: Distribution of initial desired speeds.

3 General Observations

In order to get a first understanding of a vehicular ad-hoc network's topology and its dynamics we investigated the highway scenario from the previous section in a qualitative manner. Of particular interest was the theoretical connectivity of the ad-hoc network formed by the cars. One question we wanted to answer was whether or not it is necessary to route packets over oncoming traffic in order to get acceptable connectivity. This question is important since routing over oncoming traffic implies fast topological changes and potential problems on the physical level (doppler effect, etc.). As a simplification we first assumed that any two nodes can communicate when they are no more than 250 meters apart (approximating the behavior of IEEE 802.11). With the given average density of nodes (6 per lane per kilometer) network partitioning should then be very rare if the positions of the nodes were equally distributed.

In order to determine the connectivity we followed a designated node for 200 seconds on a 10 km path. For

---

[*] One has to note that speeds and lane usage depend on national regulations.

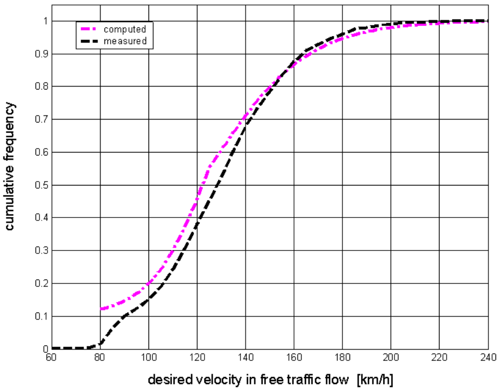

Figure 3: Computed vs. measured cumulative distribution of initial desired speeds.

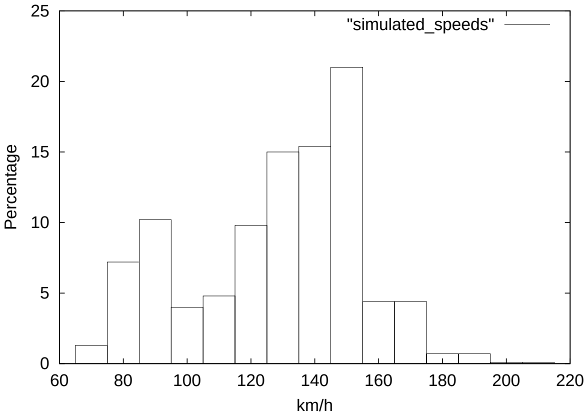

Figure 4: Distribution of speeds in the simulated scenario with 6 vehicles per kilometer and lane.

each node in a 3 km range of that node we calculated which other node could be reached directly and visualized this using Ad-Hockey. This was done twice: in the first experiment there was no communication allowed between vehicles driving in opposite directions. In the second experiment all vehicles on all four lanes were allowed to communicate with each other.

The result was converted to an MPEG video which can be downloaded from our web-server. A typical example of the connectivity when the directions are treated separately is given in Figure 5(a). This figure shows that both directions are partitioned. When investigating all 200 seconds of simulation time network partitionings are rather frequent, even though the average density of nodes is quite high. Clearly, the reason for this is that the position of vehicles is not equally distributed. This is caused by situations where one slow vehicle (e.g., a truck) overtakes another slow vehicle. In these situations connectivity will often break when oncoming traffic is not used to form the ad-hoc network.

In contrast Figure 5(b) shows the same situation when nodes on all four lanes are allowed to communicate with each other. It does show that partitionings of the network can be avoided by using oncoming traffic. An investigation of the full 200 seconds shows that most of the network partitionings can be alleviated in this fashion.

In addition we were interested in understanding how the amount of network partitions depends on the communication range. 6 shows the number of partitions on a 10 km segment with respect to the communication range of each individual node. Two graphs are given in the figure: the dotted one indicates the number of partitions when only vehicles driving in the same direction are considered for fowrwarding while the other graph describes the situation where all vehicles are taken into account. It can be seen that for the typical radio range of IEEE 802.11 (250 m) there are 7 partitions when only the vehicles driving in the same direction are taken into account. This is reduced to 2 partitions when all vehicles participate in the mobile ad-hoc network. Furthermore the graphs show that a communication range of 400 m would be desirable to completely eliminate partitioning in this scenario when all vehicles are used or 1000 m if only vehicles driving in the same direction participate.

Based on these qualitative observations it seems likely that it will be necessary to route data packets over oncoming traffic even if the density of nodes headed in the same direction is quite high. If this is not done network partitionings can be frequent and each partitioning persists for a noticeable amount of time. Therefore an adequate technology for vehicular ad-hoc networks will have to support the routing of messages over oncoming traffic.

4 Ad-Hoc Routing Strategies

Following the qualitative observations we now briefly summarize different known ad-hoc routing strategies. Two of them are presented in more detail in Sections 4.2 and 4.3. Both are then quantitatively evaluated for

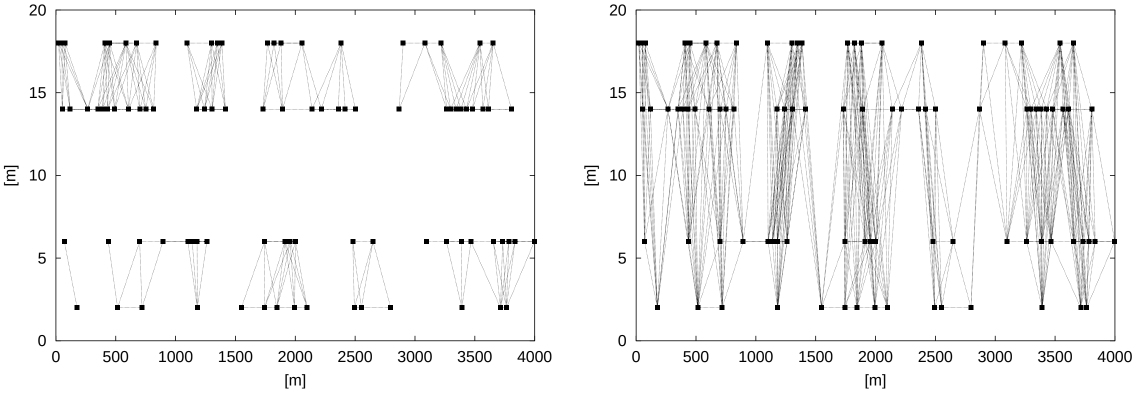

(a) Connectivity when considering only nodes headed in the same direction.

(b) Connectivity when considering nodes headed in both directions.

Figure 5: Analysis of connectivity.

their use in vehicular networks in Section 5.

4.1 Routing Protocols for mobile Ad-Hoc Networks

There are three general classes of routing protocols for mobile ad-hoc networks. *Proactive* algorithms employ classical routing strategies such as distance-vector routing (e.g., DSDV [12]) or link-state routing (e.g., OLSR [13] and TBRPF [14]). They maintain routing information about the available paths in the network even if these paths are not currently used. The main drawback of these approaches is that the maintenance of unused paths may occupy a significant part of the available bandwidth if the topology of the network changes frequently [15]. Since a network between cars is extremely dynamic we did not further investigate proactive approaches.

Reactive routing protocols such as DSR [7], TORA [16], and AODV [17] maintain only the routes that are currently in use, thereby reducing the burden on the network when only a small subset of all available routes is in use at any time. It can be expected that communication between cars will only use a very limited number of routes, therefore reactive routing seems to fit this application scenario. As a representative of the reactive approaches we have chosen DSR, since it has been shown to be superior to many other existing reactive ad-hoc routing protocols in [4].

Position-based routing algorithms require that information about the physical position of the participating nodes be available. This position is made available to the direct neighbors in form periodically transmitted beacons. A sender can request the position of a receiver by means of a location service. The routing decision at each node is then based on the destination's position contained in the packet and the position of the forwarding node's neighbors. Position-based routing does thus not require the establishment or maintenance of routes. Examples for position-based routing algorithms are face-2 [18], GPSR [8], DREAM [19] and terminodes routing [20]. As a representative of the position-based algorithms we have selected GPSR, (which is algorithmically identical to face-2), since it seems to be scalable and well suited for very dynamic networks. Examples of existing location services which map the ID of a node to its position are Homezone [21], Grid Location Service [22], and the location service part of DREAM [19]. All of these location services are proactive in the sense that they continuously communicate to maintain the position of all nodes at all times. In order to enable a fair comparison with reactive ad hoc routing strategies we have developed a trivial reactive location service which causes communication only when the position of a node is actually requested. This location service is used to determine the position of a destination in our experiments.

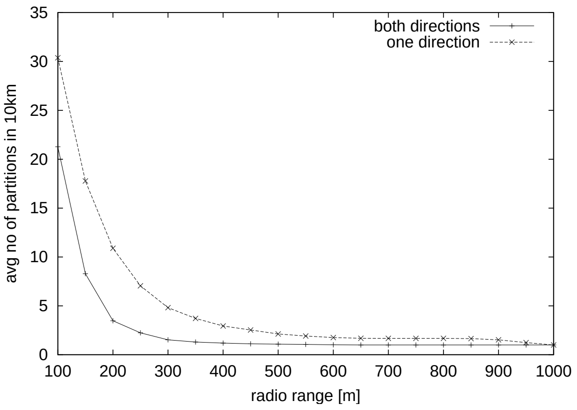

Figure 6: Number of partitions with respect to radio range

4.2 Dynamic Source Routing

Dynamic Source Routing (DSR) is typically performed in two steps: route discovery and route maintenance. A node that wants to send a packet to another node first checks its local route cache. This cache contains all valid routes the node knows about. If no route to the destination is present in the cache, a route discovery is performed. Essentially the route discovery requires that the node performing it floods the network with a *route request* that contains the ID of the node it wants to contact. Whenever the route request is forwarded by a node the forwarding node's ID is recorded in the packet. When it finally is received by the destination the route request contains a valid path from the source to the destination. The destination then sends a *route reply* back to the sender on the path contained in the route request. The sender and all nodes on the path from the destination to the sender put the route in their route cache. In order to reduce the amount of flooding, DSR employs a number of additional algorithms. For example intermediate nodes that have a valid route to the destination in their route cache may answer the route request directly with the information from the cache. In most situations these additional algorithms can prevent a full flooding of the network.

As long as two nodes communicate with each other *route maintenance* makes sure that a path between both nodes exists. When a path breaks, a packet that cannot be forwarded will generate a *route error* which is sent back to the sender of the original packet. On its way the route error causes the removal of the invalid route from the route caches of the intermediate nodes and of the sender. The sender then performs a new route discovery to find a new route to the destination. As with route discovery the route maintenance is supported by a number of additional algorithms that optimize the behavior of DSR. For a full description of DSR the reader is referred to [7].

4.3 Position Based Routing

As mentioned above, position based routing consists two building blocks: a location service and the actual forwarding of packets.

4.3.1 Reactive Location Service

Our reactive location service (RLS) is inspired by DSR route discovery: whenever the position of a node is required, the node looking for position information floods a request containing the ID of the node it is looking for. The request contains the ID and position of the requesting node. When a node receives a request with its own ID, it replies to the node looking for its position.

In order to reduce the range of the flooding an expanding ring search is performed: the flooding starts with a range of 2 hops and is repeated with a greater range when no response is received during a certain time. The range of the flooding can be increased, e.g., linearly or exponentially.

With the reactive location service there is only overhead when data actually needs to be transmitted. This makes the comparison with reactive ad-hoc routing strategies quite fair. Using one of the existing location services would produce an overhead which does not (directly) relate to the transmitted payload data. Thus any results would depend on how much payload data is transmitted.

Clearly, the overhead of the reactive location service will be generally high if communication partners are changed frequently - and thus may be inferior to exiting approaches in those situations. However, for comparison purposes it seems to be more appropriate than proactive location services. In addition the reactive location service could be optimized, using caching and prediction of a node's future location based on its speed and heading. This was not part of this study and is left for future work.

4.3.2 Greedy Perimeter Stateless Routing

In Greedy Perimeter Stateless Routing (GPSR) a node knows the position of its neighbors by means of their beacons and the position of a packets destination with the help of the location service.

With this information a node forwards incoming packets to a neighbor located in the general direction of the destination. Ideally, this process can be repeated until the destination of the packet has been reached. Unfortunately it is possible that a node represents a local optimum and has no neighbor which is closer to the destination than itself. In this situation GPSR employs an algorithm called *Perimeter Routing* which uses an algorithm for planar graph traversal to find a way out of the local optimum. The same algorithm was also proposed for face-2 in [18]. Since the topology of a vehicular network on a highway is unlikely to encounter local optima, we have turned Perimeter Routing off during our experiments.

5 Comparison of Routing Strategies

In the following we analyze the quantitative behavior of DSR and GPSR/RLS when applied to a network of vehicles.

5.1 Simulation Setup

The environment used for the simulation is based on the all-in-one distribution of ns-2.1b8a running under Linux. The `GPSR` code of Brad Karp was ported to this platform. The `DSR` code used is the one integrated in the distribution. We took a time slice of 200 seconds of the input data and a reduced kilometer range of 10 km (Position from 10 km to 20 km of the original data). This results in about 300 nodes in the scenario.

All experiments were conducted with two different MACs. One was IEEE 802.11 as provided in ns-2. The other one was an idealized MAC we implemented to abstract from MAC-specific effects. This *0-MAC* allows communication between two nodes if they are 250 meters or less apart and does not impose any upper limit on the amount of transmitted data. Collision between distinct packets that are simultaneously transmitted do not occur with the 0-MAC.

5.1.1 Communication Pattern

For the selection of the communication pattern we used the following algorithm. At any time there are 10 pairs of one sender and one receiver. These pairs are randomly selected such that they are no more than a maximum communication distance (in meters) from each other apart. In addition they are guaranteed to be theoretically able to reach each other during the time they communicate (i.e., they do not reside in different partitions). The sender then transmits 4 packets per second over a time of 5 seconds. The starting time is randomized in order to prevent synchronization. Whenever a message is successfully delivered, the receiver sends a reply. Thus we simulate typical bidirectional traffic as produced, e.g., by TCP. All packets carry a payload of 64 byte. The maximum distance between senders and receivers was varied from 500 meters to 4500 meters. Since the selection of partners is random (equally distributed) among the nodes fulfilling the constraints, sender and receiver can travel in the same or in different directions.

5.1.2 A Note about Border Effects

When simulating a linear street scenario, one has to consider border effects. For instance, a node leaving the studied area has to be deactivated, for its real position is off scope of the simulation. To accomplish that, we used the energy model of ns-2. If a node reaches the border of the simulated area, it is deactivated and reactivated (again) when it (re)-enters the scenario. In our scenario, since no node is allowed to travel backwards, each node is activated exactly and deactivated at most 1 time. Of course, when a node is deactivated, it stops sending `GPSR` beacons.

5.1.3 DSR Setup

The parameters originally set in the ns-2 implementation of `DSR` were kept for our simulation. The only modification was done to increase the maximum hop distance that a DSR route can span from 16 to 32 so that it is possible to reach all destinations even in the 4500 meter communication pattern. For a deeper understanding of `DSR` optimization, please refer to [7]. In our simulation `DSR` uses the promiscuous mode of the network interface to investigate all packets receivable regardless of the destination address.

5.1.4 GPSR and RLS Setup

Greedy Perimeter Stateless Routing was set up as follows (for a more detailed description of the parameters please refer to [8]): the beacon information of a node, i.e. its own position, is piggybacked on every packet (data packets and location service packets) that it forwards. When piggybacking a beacon the node resets the timer for the scheduling of its next beacon. We varied the beacon interval between $\{0.25, 0.5, 1, 2\}$ seconds to study its influence on the rate of successfully delivered packet and routing overhead. We make use of the MAC callback feature, enabling a node to reroute packets still buffered by the MAC if a MAC link breaks. Although this is a violation of the strict layer separation, the gain of it is remarkable according to [8].

Our Reactive Location Service was used first with linear expanding ring search and then with exponential expanding ring search. The timeout value for triggering the flooding with an increased range was set to 100ms multiplied with the number of hops in the last cycle. The maximum hop-count for the flooding was set to 32, the same value used by DSR. This should enable us to reach any node in the simulated area. Each data packet and each reply sent in response to a data packet contains the ID and location of its sender and its receiver. Thus the location information about a communication partner is updated by the receipt of a packet from that communication partner.

5.2 Simulation Results

5.2.1 0-MAC

In order to gain an impression that is unaffected by the properties of the MAC we started the simulations by using the 0-MAC. The first experiments were conducted for the position based approach with a linear expanding ring search (increase of 1 hop per cycle). Surprisingly we had many cases where a destination node was not reached by the flooding. A more detailed analysis helped us to understand the reason for this: the problem occurs when two vehicles want to communicate which drive in different directions. For the first flooding a range of 2 was used while the vehicles were n hops apart (n was greater than 2). Flooding with range 2 therefore remained without success. However, during the time required for the first cycle to time out, the cars moved in opposing directions so that they now were at least $n+1$ hops apart: the expanding ring search was slower than the vehicles. We concluded that for vehicle communication linear expanding ring search is not suitable as location service. Thus in the following we only consider exponentially expanding ring search.

One key performance metric for the suitability of a given approach is the rate of successfully delivered packets. Figure 7 shows this metric for DSR and GPSR with increasing maximum communication distances. There is just one plot for GPSR since all tested beaconing frequencies provided the same results in all ranges. This is no surprise since the flooding for the location service allows to piggy back the beacon information of all nodes between sender and receiver at the beginning of the communication. Furthermore the data packets sent will also be used for piggy backed beacons and keep the information about neighbors up-to-date. We tested beaconing frequencies with up to 16 seconds between beacons without major change in the outcome of the experiment.

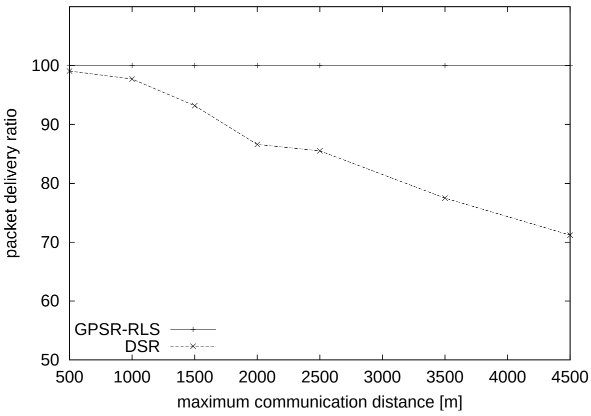

Figure 7: Packet Delivery Ratio w.r.t. Maximum Communication Distance using the 0-MAC

Figure 7 can be interpreted as follows: as expected the rate of successfully delivered packets for DSR diminishes when the maximum communication distance becomes larger. This is caused by the fact the DSR needs to maintain a route from the sender to the receiver which becomes harder when the length of the route increases. The position based approach stays at the perfect packet delivery rate of 100% for all dis-

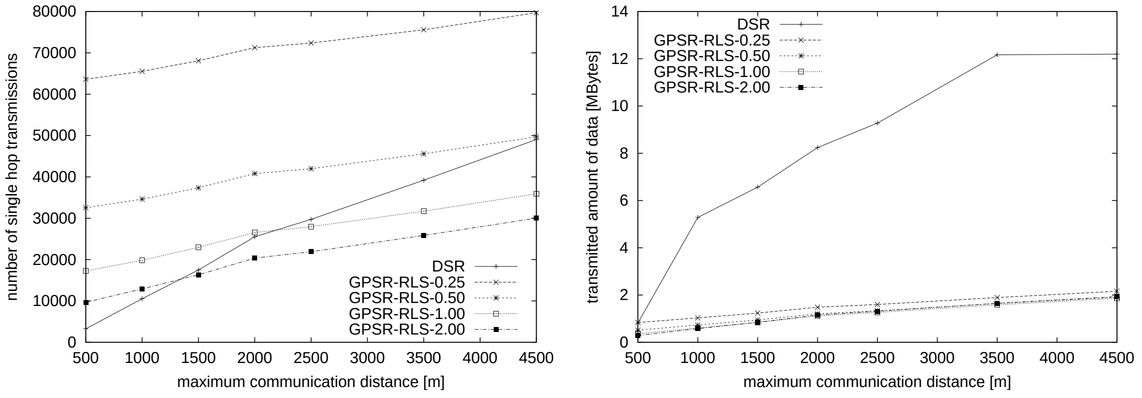

(a) Number of 1-hop packet transmissions.

(b) Volume of transmitted data.

Figure 8: Analysis of communication costs.

tances.[†] This can be explained by the properties of position based approaches: packet drops can occur only for one of the following three reasons: (1) if a local maximum is reached. This is extremely unlikely in our scenario. (2) If the information about the position of the local neighbors is inaccurate. Again this is very unlikely since the flooding of the location service in combination with piggy backed beacons will provide nearly perfect information about the neighbors. (3) If the information about the position of the destination is inaccurate. This is also very rare, since using the 0-MAC the reply containing the position of the destination requires only minimal time to reach the sender, thus it is very accurate when the data packet is transmitted.

Besides looking at the delivery rate it is also important to investigate how many packets and how much data is required to transmit a certain amount of payload data. We therefore measured the total number of one hop transmissions that occurred over the whole lifetime of the simulation. This is shown in Figure 8(a). Both unicast and broadcast messages are included in this figure. For GPSR we show the communication costs for all beaconing frequencies. It can be seen that the value for DSR starts low when the maximum communication distance is small and grows fast with increasing communication distances. This is caused by the increase in overhead for route establishment and maintenance which are the main sources of packets for DSR (besides the actual data packets). GPSR/RLS on the other hand starts at a higher value and then increases more slowly. Furthermore it can be noticed that the communication overhead scales almost linearly with the beaconing frequency. The reason for this behavior is that beacons are the dominating source of one hop transmissions in position-based routing. These are independent of the maximum distance between communication partners. Since the packet delivery ratio is almost independent of the beaconing frequency and since beaconing provides the dominating amount of one hop packet transmission it seems appropriate to use a fairly low beaconing frequency when employing GPSR/RLS for vehicular networks.

Figure 8(b) displays the total amount of data used in form of single hop transmissions. It demonstrates that DSR needs significantly more data than GPSR for all examined maximum communication range values and beaconing frequencies. This is caused by the size of the packets needed to establish and maintain routes in DSR. Since these packets need to carry a source route from the sender to the receiver they can become quite large. As a contrast the packet size of GPSR/RLS is very small: all that is required is the position information and ID of the sender (and of the receiver if it is a data packet).

[†] It should be emphasized that we did not try to “optimize” the simulation to achieve this figure. In fact we would have preferred a somewhat less perfect result. We invite people to validate these results and will therefore put up everything required to run the simulation on the web.

5.2.2 IEEE 802.11

In a second round we repeated the experiments using the default implementation of IEEE 802.11 in ns-2 as MAC. Given the results from the previous section we expected similar but somewhat less optimal results. In our initial experiments with IEEE 802.11 we were surprised to see that GPSR/RLS actually performed similar and sometimes worse than DSR in respect to the rate of successfully delivered packets. In particular the exponential expanding ring search frequently failed to reach the destination node. Investigating this problem we noticed that the flooded packets tended to synchronize themselves such that they cause collisions at the MAC layer. In IEEE 802.11 broadcast packets that are affected by such a collision are not retransmitted and remain lost. Thus the synchronized broadcasting of packets can lead to a complete whipeout of the affected packet. As a consequence we introduced a jitter when sending broadcast packets for the expanding ring search. This solved the problem.

Figure 9 shows the packet delivery rate for the simulation with IEEE 802.11. Generally the outcome is very similar to the 0-MAC case. However, there is one minor detail that is worth mentioning: for GPSR/RLS we had some runs where data packets got lost, even though the vast majority of runs did complete without a single packet loss. The main reason for those losses was that beacons and broadcast packets from the location service would still sometimes collide. Thus the information about the position and availability of neighbors is less accurate in the simulation runs with IEEE 802.11. This sometimes causes a forwarding node to be ignorant of the only neighbor with forward progress in the direction of the destination. The cost for the communication remained very similar to that of the 0-MAC case and is therefore not shown here.

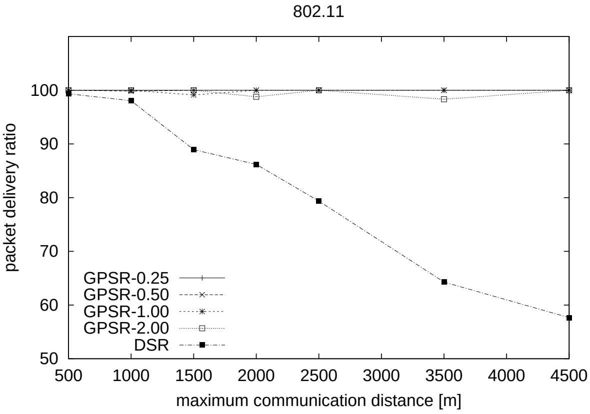

Figure 9: Packet Delivery Ratio w.r.t. Maximum Communication Distance using IEEE 802.11

6 Conclusions and Outlook

In the near future communication between vehicles will increase the safety and the comfort of passengers. In order to remain independent of a potentially very expensive infrastructure, routing approaches for mobile ad-hoc networks can be used. In this paper we demonstrated that currently position-based approaches seem to be particular promising for communication in vehicular networks. They provide a very high rate of successfully delivered packets even over many hops. Furthermore their overhead in terms of one hop transmissions and transmitted bytes scales better than that of reactive approaches. The reason for this is that position-based routing does not have to maintain routes and instead performs forwarding ‘on the fly’.

However, research in this area is far from being complete. For example, reactive approaches such as DSR might be improved by considering the movement of the individual nodes in the routing decision. This way they could give preference to routes over vehicles driving in the same direction and thus minimizing the number of topological changes that might lead to link breaks. The trivial location service used for position based adhoc routing could be significantly optimized by using caching and prediction of a node’s future location, based on its speed and heading.

Furthermore there are many parameters in both reactive and location based ad-hoc routing that could be tuned to optimize the performance for vehicular communication. These include single hop radio transmission range, beaconing frequency, optimized flooding strategies and many more. Finally it will be interesting to investigate and compare the behavior of the routing strategies when used in a scenario involving city traffic.

References

[1] Robert Morris , John Janotti , Frans Kaashoek , Jinyang Li , and Douglas S.J. DeCouto, “Carnet: A scalable ad hoc wireless network sys-

tem," in *Proceedings of the 9th ACM SIGOPS European workshop: Beyond the PC: New Challenges for the Operating System*, Kolding, Denmark, September 2000, p. 127ff.

[2] "The FleetNet project," http://www.fleetnet.de.

[3] Jinyang Li, Charles Blake, Douglas S.J. De Couto, Hu Imm Lee, and Robert Morris, "Capacity of ad hoc wireless networks," in *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking (MobiCom 2001)*, Rome, Italy, July 2001, pp. 61–69.

[4] Josh Broch , David A. Maltz , David B. Johnson , Yih-Chun Hu , and Jorjeta 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, U.S.A., October 1998, pp. 85 – 97.

[5] "The ns-2 network simulator," http://www.isi.-edu/nsnam/ns/.

[6] "The cmu monarch wireless and mobility extensions to ns-2," http://www.monarch.cs.ċmu.edu/-cmu-ns.html.

[7] David B. Johnson and David A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing*, Tomasz Imielinske and Hank Korth, Eds., vol. 353. Kluwer Academic Publishers, 1996.

[8] Brad Karp and H. T. Kung, "GPSR: Greedy perimeter stateless routing for wireless networks," in *Proceedings of the 6th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*, Boston, MA, U.S.A., August 2000, pp. 243–254.

[9] D. Helbing, "Traffic and related self-driven many-particle systems," *Rev. Modern Physics*, vol. 73, pp. 1067–1141, 2001.

[10] D. Vollmer, B. Balasubramanian, and E. Siegert, "Fahrtsimulation unter realistischen Umfeldbedingungen (in German)," VDI-Berichte, 1992.

[11] T. Benz, L. Schäfers, C. Stiller, and D. Vollmer, "Feasibility study on truck planning on european motorways," Deliverable D08.1 of ITS project PROMOTE-CHAUFFEUR, 1999.

[12] Charles E. Perkins and Pravin Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (DSDV)," in *Proceedings of ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, 1994.

[13] Thomas Clausen , Philippe Jacquet , Anis Laouiti , Pascale Minet , Paul Muhlethaler , Amir Quayyum , and Laurent Viennot , "Optimized link state routing protocol," Internet Draft, draft-ietf-manet-olsr-05.txt, work in progress, October 2001.

[14] Richard G. Ogier , Fred L. Templin , Bhargav Bellur , and Mark G. Lewis , "Topology broadcast based on reverse-path forwarding (tbrpf)," Internet Draft, draft-ietf-manet-tbrpf-03.txt, work in progress, November 2001.

[15] S. R. Das, R. Castaneda, and J. Yan, "Simulation based performance evaluation of mobile, ad hoc network routing protocols," *ACM/Baltzer Mobile Networks and Applications (MONET) Journal*, pp. 179–189, July 2000.

[16] Vincent D. Park and M. Scott Corson, "A highly adaptive distributed routing algorithm for mobile wireless networks," in *Proceedings of IEEE INFOCOMM*, 1997, pp. 1405–1413.

[17] Charles E. Perkins and Elizabeth M. Royer, "Ad-hoc on-demand distance vector routing," in *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, February 1999, pp. 1405–1413.

[18] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia, "Routing with guaranteed delivery in ad hoc wireless networks," in *Proc. of 3rd ACM Intl. Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications DIAL M99*, 1999, pp. 48–55.

[19] Stefano Basagni, Imrich Chlamtac, Violet R. Syrotiuk, and Barry A. Woodward, "A distance routing effect algorithm for mobility (dream)," in *ACM MOBICOM '98*. ACM, 1998, pp. 76 – 84.

[20] Ljubica Blazevic , Silvia Giordano , and Jean-Yves Le Boudec , "Self-organizing wide-area

routing," in *Proceedings of SCI 2000/ISAS 2000*, Orlando, July 2000.

[21] S. Giordano and M. Hamdi, "Mobility management: The virtual home region," Tech. Rep., October 1999.

[22] J. Li, J. Jannotti, D. S. J. De Couto, D. R. Karger, and R. Morris, "A scalable location service for geographic ad hoc routing," in *Proc. of the 6th Annual ACM/IEEE Int. Conf. on Mobile Computing and Networking (MOBICOM) 2000*, Boston, MA, USA, 2000, pp. 120–130.