Overcoming Fragmentation in Mobile Ad Hoc Networks

Linda Briesemeister and Günter Hommel

Abstract: **We present an approach to multicast messages among highly mobile hosts in ad hoc networks. We suggest a new definition of a multicast that suits the special needs of inter-vehicle communication: rather than explicit identification, a multicast group is defined implicitly by location, speed, driving direction and time. As an example, we study a road accident that is reported to nearby vehicles. We focus on sparse deployment of the system which is likely to occur soon after the system is introduced to the market. In this state, the resulting ad hoc network tends to be disconnected. We tailor the proposed algorithm to overcome this problem of network fragmentation. Simulations show us the quality of the proposed protocol by measuring how many vehicles inside a multicast area are informed in time under various conditions.**

Index Terms: **Inter-vehicle communication, mobile ad hoc network, multicast, global positioning system.**

I. INTRODUCTION

Recently, mobile computing has become a hot topic in research. Although computer and communication devices are becoming smaller and more powerful, mobility still challenges applications of mobile computing especially in the area of ad hoc networking. A mobile ad hoc network consists of mobile hosts that communicate via wireless links. Due to mobility, the topology of the network changes continuously and wireless links break down and reestablish frequently. Moreover, an ad hoc network operates in the absence of fixed infrastructure forcing the hosts to organize the exchange of information decentrally.

A prominent application of mobile ad hoc networks is direct wireless communication between vehicles in road traffic. In this application, the vehicles are equipped with a computer controlled radio modem allowing them to contact other equipped vehicles in their vicinity. Adhering to the abstract definition of an ad hoc network, we assume no fixed infrastructure to support the communication. Still, some aspects of it make inter-vehicle communication distinct from ad hoc networking in general—namely the high mobility and the anonymity of hosts and as a consequence of both, a modified definition of multicast.

We believe that the best applications of inter-vehicle communication are to provide improved comfort and additional safety in driving. Our aim is to make these applications feasible by enabling the dissemination of information among participating vehicles. In contrast to applications in cooperative driving [1], [2], platooning and automated highways [3]–[6], we are able to relax the requirements of high bandwidth, expensive equipment and infrastructure, and most importantly 100% deployment. Tests with relatively cheap off-the-shelf devices [7], [8] have shown the general feasibility of radio modems in the 2.4 GHz industrial, scientific, and medical (ISM) band. As an example application, an equipped vehicle identifies itself as crashed by vehicular sensors that detect events like airbag ignition. Then, it can report the accident instantly to equipped vehicles nearby. We present an algorithm to disseminate such a message among the other equipped vehicles on the road. Multihopping allows us to enlarge the area in which a vehicle could receive the message. If the message reaches a vehicle for which the warning is relevant, the driver can be informed early by the system. Thus, we intend to help the driver cope with a potentially dangerous or inconvenient situation.

Manuscript received mm dd, year.

L. Briesemeister is with DaimlerChrysler Research and Technology, Alt-Moabit 96a, 10559 Berlin, Germany, Linda.Briesemeister@DaimlerChrysler.c-om.

G. Hommel is with the Department of Computer Science, Technical University of Berlin, Germany.

II. RELATED WORK

Many studies in ad hoc networking [9]–[13] propose mobility patterns in the two-dimensional plane. There, the hosts change their speed and direction more or less randomly. However, vehicles in road traffic typically follow the road which allows us to reduce mobility to one dimension. Furthermore, vehicles on a highway often drive at 130 km/h and more which is much faster than the papers [9], [10] presume. Hong *et al*. [14] have shown the impact of mobility patterns on performance measures of ad hoc network protocols. Thus, we study the proposed protocol under the conditions of a highway traffic model.

Another impediment in mapping inter-vehicle communication to previous research in ad hoc networking is the anonymity of participating hosts. As an essential requirement in networking, every host vehicle possesses a unique identifier. But the set of existing identifiers can easily exceed a practical size of fixed host or server tables. Plus, newly manufactured vehicles equipped with the system join the set of existing identifiers whereas the identifiers of crashed or scrapped vehicles leave the set. In contrast to the huge set of possible identifiers, an ad hoc network formed in reality on the road will only connect a small subset of identifiers. In applications of inter-vehicle communication, a vehicle often needs to address other nearby vehicles whose identities are a priori unknown rather than sending a packet to a specifically identified vehicle.

In our sample scenario with a road accident, the crashed vehicle wants to inform the other vehicles that are approaching the hazardous area. To make the system work, the vehicles need to be aware of their locations. Many vehicles do or will soon utilize navigation systems like the Global Positioning System (GPS). Although today's GPS receivers are accurate to within 100 m, we expect dramatic improvement during the next several years. Future navigation systems will use differential correction or integrate inertial sensors to enhance the accuracy of positioning down to a few meters or better [15], [16]. Assuming that equipped vehicles know their location more or less accurately, they can direct messages to a specific geographical area. Imielinski and Navas [17] and Ko and Vaidya [11] proposed the idea of using geographic constraints to specify the destination of a packet to be routed and coined the term "geocast."

However, in our application, aspects other than location determine whether a vehicle belongs to the multicast group. Taking the driving direction into account, a vehicle can distinguish more reliably whether it is approaching the dangerous spot or not. Also, if it employs a digital road map, the vehicle improves its ability to classify the situation. On divided highways, an accident usually does not harm vehicles in the other driving direction. Finally, the velocity of a vehicle puts an individual deadline on message delivery for each potential recipient inside the multicast region. When the vehicle receives the warning after it has passed the accident, the message is useless to the driver. Due to the high speed of mobile nodes, the ultimate point where a vehicle has to be informed is within braking distance of the accident. We use (1) to calculate each vehicle's braking distance depending on its velocity. Hence, the multicast group is further limited to those vehicles inside the multicast region that are still able to stop in front of the accident. We define the multicast group implicitly by applying these constraints on individual vehicles.

$$dist_{brake}(v) = v \cdot \Delta t_{reaction} + \frac{v^2}{2 \cdot b_{max}} \tag{1}$$

where $\Delta t_{reaction}$: reaction time of driver $= 1$ sec

b_{max} : maximum deceleration $= 4.4$ m/sec^2

III. PROPOSED ALGORITHM

In mobile ad hoc networks, three basic categories of multicast algorithms exist. The pro-active approach precomputes paths to all possible destinations and stores this information in routing tables. To maintain an up-to-date database, routing information is periodically distributed throughout the network. As mentioned in the previous section, the huge set of possible hosts plus the high mobility makes this approach impractical for inter-vehicle communication.

As a second category, routing and multicast protocols create paths to other hosts on demand. The idea is based on a query-response mechanism and is sometimes called "reactive multicast." In the query phase, a node explores the environment. Once the query reaches the destination, the response phase is issued and establishes the path. Still, the sender requesting a route to a destination must know its identity. If this knowledge is missing, the sender first has to collect information about the network topology.

In the third category, multicasting algorithms are simply flooding the network. Every node receiving a message forwards it to a list of neighbors. Flooding a network acts like a chain reaction that can result in exponential growth. Nonetheless, it has the advantage of working with no knowledge about the underlying network topology. Thus, the query phase in reactive multicasting often uses flooding to find a path.

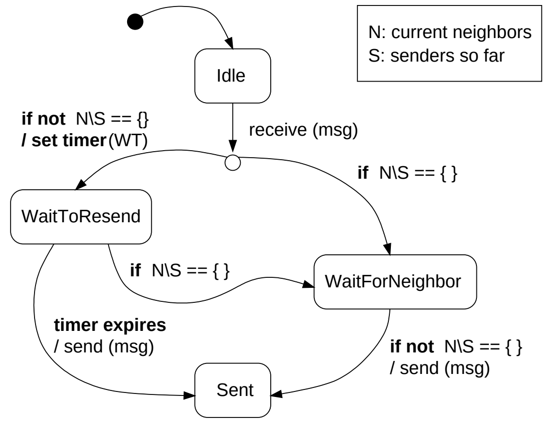

Fig. 1. State transition diagram of proposed algorithm.

Assuming that warning about an accident is a relatively rare event, it seems suitable to design the protocol working on demand. However, a warning message is not much bigger than a control packet sent during a request phase. Another drawback of reactive multicast schemes is the explosion of ACKs sent to control the establishment of paths. Hence, we avoid the response phase and simply apply flooding to reach the multicast group. We also take advantage of the broadcasting nature of radio waves; a packet sent by one host can reach multiple receivers simultaneously. Thus, the number of sending activities only increases linearly with the number of hosts although the number of packets received still grows exponentially.

The Location-Based Multicast (LBM) protocol described by Ko and Vaidya [11] also uses flooding and is therefore similar to our approach. The region to which the geocast should be delivered is named the "multicast region." A "forwarding zone" contains the multicast region and connects it to the source node. LBM differs from our algorithm in that LBM limits the flooding process to nodes inside the forwarding zone whereas in our scheme potentially everybody participates in the dissemination process as long as messages do not exceed a maximum number of hops.

In previous studies on flooding [18], we encountered the problem of a fragmented network due to a small number of equipped vehicles on the road. Obviously, the success of the system in the market place depends on the system functioning and producing visible results with only a few vehicles being equipped. Furthermore, our simulations in [18] proved that reaching the maximum of addressed vehicles only takes one second covering a stretch up to 5 km. Realizing that fast delivery is not a crucial factor, we propose the idea of allowing nodes to not forward the message until new receivers move into their vicinity.

We assume our protocol works on top of a data link layer (DLL) that keeps track of the neighboring nodes. Thus, the system is notified when new neighbors enter the vehicles' transmission area or existing neighbors move out of range. The crashed vehicle starts sending the warning message when it senses a neighboring vehicle. After initiation of the message dissemination, all other equipped vehicles perform the following algorithm. Fig. 1 depicts a state transition diagram of this algorithm.

Each vehicle maintains the set N of neighbors. It constantly updates N according to the notification from the data link layer. Also, each vehicle associates the set S with the warning message. Every time the system receives the message from a sender, it adds the sender's identity to S. On the first reception of the message, S is initialized with the corresponding sender identity and the system switches into one of the two states "WaitToResend" or "WaitForNeighbor."

If $N \setminus S \neq \emptyset$, the system has neighbors other than the sender of the previously received message and it enters the "WaitToResend" mode. We assume that the message header contains the position of its sender. By knowing its own position, the system determines a waiting time WT depending on the distance d to the sender such that the waiting time is shorter for more distant receivers as shown in (2):

$$WT(d) = -\frac{MaxWT}{Range} \cdot \hat{d} + MaxWT \tag{2}$$

$$\hat{d} = \min\{d, Range\}$$

$$\text{where} \quad MaxWT : \text{maximum waiting time}$$

$$Range : \text{transmission range.}$$

To understand our motivation to wait rather than to resend the message immediately, consider the broadcasting nature of radio waves. Multiple hosts can receive the same packet simultaneously. Then, an immediate resending would cause burst-like traffic on the channel. The medium access control (MAC) in ad hoc networks is often based on a carrier sense multiple access (CSMA) mechanism. It is well known that CSMA suffers from instability when the capacity of the channel is reached [19], [20]. Hence, we try to avoid peak load by forcing the receivers to wait. Using the function WT, mainly hosts at the border of the reception area take part in forwarding the message quickly.

While the system awaits the moment to resend, it still updates the sets N and S. If on any of these updates the condition $N \setminus S \neq \emptyset$ does not hold anymore, the system switches into the "WaitForNeighbor" state. Otherwise, it forwards the message after the calculated waiting time is over.

If $N \setminus S = \emptyset$, then there are no new receivers nearby and the system switches into "WaitForNeighbor" mode. In this state, the system waits until an update of N occurs such that $N \setminus S \neq \emptyset$. Then, the system forwards the message.

IV. SIMULATION MODEL

To study our approach to multicasting, we implemented the following model of an inter-vehicle communication scenario in SHIFT [21] which was developed within the PATH program at the University of California, Berkeley.

We assume that the vehicles use omnidirectional antennas implying that a sender can transmit to multiple hosts simultaneously. When they are within transmission range of a sending vehicle, all other equipped vehicles potentially receive the data packet. Having tested two 2.4 GHz radio modems [7], we measured approximately 600 m as the maximum distance for receiving data. We choose the maximum waiting time until a vehicle forwards a packet to be 40 ms. The dissemination of the message is also ultimately controlled by the maximum hops it can take. We limit the propagation of a message to 20 hops.

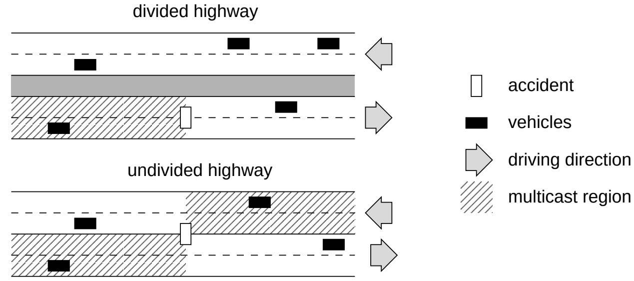

Fig. 2. Sample scenarios for different road types.

As an example application of our approach, we demonstrate the equipped vehicles distributing a warning message about an accident in road traffic. We model a straight road 10 km long with two lanes in each direction. The accident happens in the middle of the simulated stretch. Two different road types are considered: a divided highway and an undivided highway.

When measuring the performance of our multicast protocol, the multicast region in which drivers should be informed about the accident depends on those road types. For the road model of the divided highway, the multicast region covers the area behind the accident on the side of the highway where the accident happens. On the undivided highway, the vehicle having an accident can affect both driving directions. Hence, all vehicles approaching the position of the accident are part of the multicast region. Refer to Fig. 2 that sketches the scenarios described above.

Each vehicle on the road moves at a constant, randomly chosen velocity. For the sake of simplicity, we do not model complex maneuvers like lane changes or overtaking. Furthermore, we assume that the traffic is relatively dense but is still free flowing. We then determine the distribution of velocity from a traffic model by Heidemann [22]. Two parameters define this model: the average velocity and the traffic density. For our highway scenario, the velocity varies around the value of 130 km/h and the average traffic density is 5 vehicles per km per lane.

The process of message dissemination depends heavily on the number of equipped vehicles on the road. When the system is introduced to the market, only a small number of vehicles will be equipped. We designed our algorithm to overcome the problem of fragmentation in sparsely connected networks. Thus, we incremented the percentage of equipped vehicles from 1% up to 10 percent. To study the transition to full market share, we also considered values of 15, 20, 25, 50, and 100% deployment. We executed 100 simulation runs for each set of parameters.

V. SIMULATION RESULTS

A. Metrics

The simulation starts by initiating the message dissemination from the crashed vehicle in the middle of the simulated road stretch once. At this moment, all equipped vehicles inside the multicast region determine their braking distance according to (1). The multicast group consists of all vehicles in the multicast region that are still able to stop in front of the accident according to their braking distance. The simulation proceeds until all members of the multicast group have passed the accident which

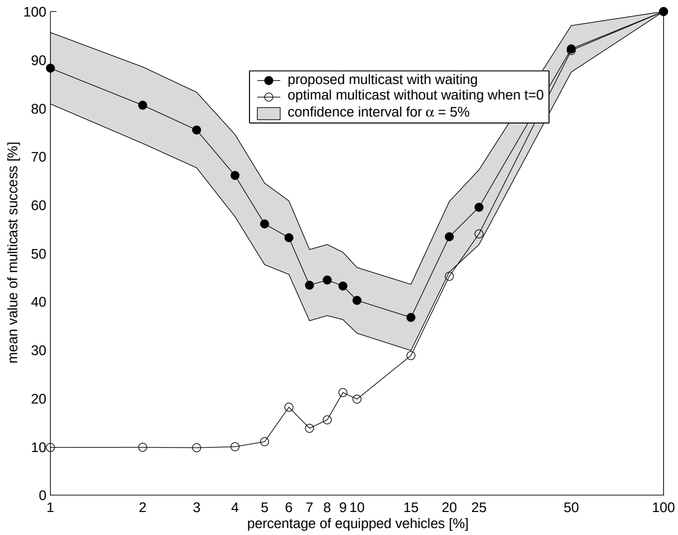

Fig. 3. Results for success of multicast on divided highway.

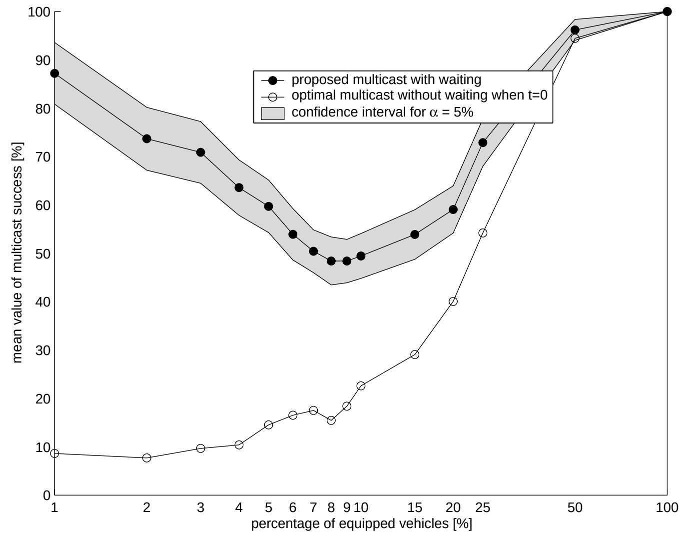

Fig. 4. Results for success of multicast on undivided highway.

denotes a stable state regarding our metrics. Once the distance of a vehicle to the accident becomes less than the vehicle's braking distance, we note whether the vehicle is informed or not. We characterize a simulation run by the success of the multicast which is the ratio of informed vehicles to the size of the multicast group. This metric is similar to the "Accuracy of Multicast Delivery" in [11]. In the case that the multicast group has no members, the result of the success metric is undefined. Hence, we eliminate all simulation runs as invalid that meet this criteria. The mean values over all valid simulation runs for different percentages of equipped vehicles on the road are shown in Figs. 3 and 4. Note, that we scaled the horizontal axis logarithmically to enhance visibility of the results for small deployment of 10% and less.

We compare the success of our protocol to the success of an optimal but unknown multicast protocol without waiting when $t = 0$. For simplicity, the time $t = 0$ represents the moment the accident happens. Then, we take a snapshot of the network topology resulting in a graph $G = (V, E)$. The equipped vehicles form the set of nodes V. An edge between two nodes denotes that the corresponding vehicles are within transmission range of each other. Let R be the set of equipped vehicles that are reachable from the crashed vehicle in this topology i.e., a path from the crashed vehicle to the equipped vehicle exists in graph G. If the set M denotes the multicast group, the set $R \cap M$ is a maximum subset of the multicast group that an optimal algorithm could inform. Note, that it is not necessarily true that $R \subset M$ or $M \subset R$. We say that $\frac{|R \cap M|}{|M|}$ is the success of an optimal multicast without waiting when $t = 0$. Again, if $M = \emptyset$ the result of the success metric is undefined. To compare our algorithm with the above sketched optimal multicast without waiting, we calculate the mean values of success on top of the reduced data set. Thus, the simulation runs with an empty multicast group have already been removed. The results for an optimal multicast without waiting when $t = 0$ are added to the curves printed in Figs. 3 and 4.

Furthermore, we determine the quality of the simulation output using independent replications analysis [23], [24]. We stopped every simulation run when a certain criterion was met. Thus, these runs fall into the category of terminating simulations. For each run, the pseudo random number generator starts with a new initial seed and hence produces a replication independent from other runs. For if the number of independent replications is large enough, a central limit theorem allows us to assume that the replicate outputs are approximately i.i.d. normal. We then compute the half-width of the $100(1-\alpha)\%$ confidence interval as stated in (3) where $t(d, p)$ represents the p quantile of Student's t distribution with d degrees of freedom. Note, that invalid simulation runs had to be discarded and the actual numbers of independent replications n can be smaller than 100. We set the confidence level to 95% corresponding to $\alpha = 0.05$. The confidence intervals for our simulation results are marked gray in Figs. 3 and 4. Refer to Table 1 for numerical values of the confidence intervals.

$$M_p \pm t(n-1, 1-\frac{\alpha}{2}) \cdot \sqrt{\frac{V_p}{n}} \tag{3}$$

Table 1. Results with confidence intervals for $\alpha = 5\%$.

percentage of equipped vehicles [%]	success on divided highway [%]	success on undivided highway [%]
1	88.27 ± 7.39	87.22 ± 6.39
2	80.64 ± 7.92	73.69 ± 6.50
3	75.50 ± 7.82	70.87 ± 6.37
4	66.09 ± 8.49	63.59 ± 5.73
5	56.08 ± 8.42	59.71 ± 5.43
6	53.24 ± 7.61	53.93 ± 5.31
7	43.42 ± 7.37	50.42 ± 4.43
8	44.48 ± 7.35	48.42 ± 4.96
9	43.25 ± 6.96	48.41 ± 4.48
10	40.27 ± 6.80	49.46 ± 4.64
15	36.74 ± 6.84	53.90 ± 5.13
20	53.44 ± 7.30	59.07 ± 4.87
25	59.53 ± 7.72	72.90 ± 4.90
50	92.29 ± 4.81	96.20 ± 2.17
100	100.00 ± 0	100.00 ± 0

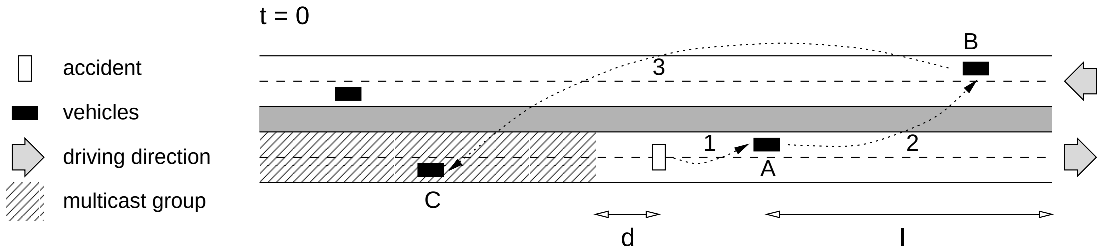

Fig. 5. Example scenario with only a few equipped vehicles.

where M_p : mean of data with parameter p
V_p : standard deviation of data with parameter p
n : number of independent replications

B. Interpretation of Results

The results on both road types show the same characteristics. When only 1% of the vehicles are equipped, the success rate reaches high values of 88.3% on the divided highway and 87.3% on the undivided highway. However, this is due to the rarity of having more than one equipped vehicle on the road. In the event of exactly one equipped vehicle on the road, our algorithm always yields 100% success rate because the crashed vehicle waits to initiate the message dissemination until it senses a neighbor. The advantage of using the ability to detect neighbors in our algorithm is clearly indicated by the immense difference in the values of an optimal multicast without waiting when $t = 0$. We achieve substantially better results than an optimal multicast algorithm would reach for deployment up to 25 percent.

In the beginning, the curves of the proposed protocol for both road types first descend for higher deployment of the system. We explain this behavior by looking at the example situation depicted in Fig. 5. As soon as the traffic model consists of a few equipped vehicles, it becomes essential who is informed first. The crashed vehicle only initiates the message once. Thus, if the receiver A of the first packet happens to be driving away from the accident, it transports the message only to the remaining stretch l of the road. Then, the chance of meeting an oncoming receiver B that is fast enough to carry the message back to vehicle C before C gets into braking distance d of the accident is small. Hence, the performance on multicast success can decrease even though the percentage of equipment increases. The lowest average success for the divided highway is 36.7% for 15% equipped vehicles and 48.4% for 9% equipped vehicles on the undivided highway. These figures correspond to a mean multicast group size of 7.0 and 8.7 vehicles respectively. Note, that the total number of equipped vehicles on the road are approximately four times the size of the multicast group on the divided highway and twice the multicast group size on the undivided highway. Further research on the group size is necessary to understand if this is a critical value.

Above deployment of 15%, the success increases with higher percentages. As expected, the significant advantage of waiting for neighbors vanishes as soon as the network becomes less fragmented. Then, the algorithm acts like pure flooding and reaches a sufficient number of destinations in time.

VI. CONCLUSION AND OUTLOOK

We presented an approach to multicast a message among highly mobile hosts like vehicles in road traffic. In contrast to classic multicasting which requires a fixed set of addresses, in this approach the role of the vehicles determine the destination of the message implicitly. As an example, vehicles equipped with inter-vehicle radio communication propagate a warning message about an accident. The proposed and implemented protocol suits especially sparsely connected networks due to small deployment soon after the system is introduced to the market. Also, we modeled two different road types: a divided and an undivided highway.

Simulation results are generated for different percentages of equipped vehicles on the road. When compared to an optimal but unknown multicast algorithm that floods the network instantly, we achieve significantly better results with our approach to waiting for neighbors. When less than 10% of the vehicles are equipped, our protocol exceeds any instant multicast algorithm at least by 20.4 percent. For all investigated deployment rates, our proposed algorithm reaches no less than 36.7% (divided highway) and 48.4% (undivided highway) of the multicast group. In turn, the compared optimal multicast algorithm resides at first below 20% and arrives at a global minimum of only 9.8% (divided highway) and 7.7% (undivided highway) due to the fragmentation of the network. Thus, our algorithm indeed overcomes the problem of a sparsely connected network by taking advantage of the high velocity and the mobility pattern of hosts in inter-vehicle communication.

Future research will focus on the behavior of multiple initiated messages. Also, the overhead in terms of unnecessary packets being sent and memory requirements of our protocol has to be studied. Even though we believe that the underlying protocol layers MAC and DLL are feasible, we wish to incorporate them into a more detailed simulation program.

REFERENCES

[1] W. Kremer and W. Kremer, "Vehicle density and communication load estimation in mobile radio local area networks (MR-LANs)," *42nd Vehicular Technology Society Conf.*, vol. 2, pp. 698–704. May 1992.
[2] M. Aoki and H. Fujii, "Inter-vehicle communication: Technical issues on vehicle control application," *IEEE Commun. Mag.*, vol. 34, no. 10, pp. 90–93, Oct. 1996.
[3] S. R. Sachs and P. Varaiya, "A communication system for the control of automated vehicles," PATH Technical Memorandum 93-5, University of California, Berkeley, Sept. 1993.
[4] W.-Y. Li, "Design and implementation of digital radio communications link for platoon control experiments," PATH Research Report UCB-ITS-PRR-95-2, University of California, Berkeley, Jan. 1995.

[5] O. Gehring and H. Fritz, “Lateral control concepts for truck platooning in the CHAUFFEUR project,” *4th World Congress on Intelligent Transport Systems*, Oct. 1997.

[6] M. Ohtomo *et al.*, “Automatic following system utilizing vehicle-to-vehicle communication,” *IEEE International Conf. on Intelligent Vehicles*, pp. 381–384, Oct. 1998.

[7] L. Briesemeister *et al.*, “Spread-spectrum funkmodems für den einsatz zur direkten digitalen fahrzeug-fahrzeug kommunikation,” Tech. Rep. FT3/AS-98-002, Daimler-Benz AG, 1998.

[8] L. B. Michael *et al.*, “DS/SS inter-vehicle communication experiments in 2.4 GHz ISM band,” *IEEE International Conf. on Intelligent Vehicles*, pp. 397–401, Oct. 1998.

[9] J. Broch *et al.* “A performance comparison of multi-hop wireless ad hoc network routing protocols,” *4th ACM/IEEE International Conf. on Mobile Computing and Networking*, pp. 85–97, Oct. 1998.

[10] C.-C. Chiang, M. Gerla, and L. Zhang, “Forwarding group multicast protocol (FGMP) for multihop, mobile wireless networks,” *ACM-Baltzer Journal of Cluster Computing: Special Issue on Mobile Computing*, vol. 1, no. 2, 1998.

[11] Y.-B. Ko and N. H. Vaidya, “Geocasting in mobile ad hoc networks: Location based multicast algorithms,” in *Proc. IEEE Workshop Mobile Computing Systems and Applications*, Feb. 1999.

[12] A. B. McDonald and T. Znati, “A mobility-based framework for adaptive clustering in wireless ad-hoc networks,” *IEEE J. Select. Areas Commun.*, vol. 17, no. 8, Aug. 1999.

[13] D. Câmara and A. A. F. Loureiro, “A novel routing algorithm for ad hoc networks,” *33rd Hawaii International Conf. on System Sciences*, Maui, Hawaii, USA, Jan. 2000.

[14] X. Hong *et al.*, “A group mobility model for ad hoc wireless networks,” *ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, pp. 53–60, Aug. 1999.

[15] A. Mathes and W. Söhne, “GPS real-time positioning and attitude determination with centimeter accuracy as a key to high-precision land vehicle control navigation,” *3rd Conf. on Land Vehicle Navigation*, pp. 93–102, June 1994.

[16] C. K. H. Wilson, S. Rogers, and S. Weisenburger, “The potential of precision maps in intelligent vehicles,” in *IEEE International Conf. on Intelligent Vehicles*, pp. 419–422, Oct. 1998.

[17] T. Imielinski and J. C. Navas, “GPS-based geographic addressing, routing, and resource discovery,” *Communications of the ACM*, vol. 42, no. 4, pp. 86–92, Apr. 1999.

[18] L. Briesemeister, L. Schäfers, and G. Hommel, “Disseminating messages among highly mobile hosts based on inter-vehicle communication,” *IEEE International Symposium on Intelligent Vehicle*, in press.

[19] L. Kleinrock and F. A. Tobagi, “Packet switching in radio channels: Part I carrier sense multiple-access modes and their throughput-delay characteristics,” *IEEE Trans. Commun.*, vol. 23, no. 12, pp. 1400–1416, Dec. 1975.

[20] F. A. Tobagi and L. Kleinrock, “Packet switching in radio channels: Part II the hidden terminal problem in carrier sense multiple-access and the busy-tone solution,” *IEEE Trans. Commun.*, vol. 23, no. 12, pp. 1417–1433, Dec. 1975.

[21] “SHIFT—The hybrid system simulation programming language,” available at http://www.path.berkeley.edu/shift/.

[22] D. Heidemann, “Ein Mathematisches Modell des Verkehrsflusses,” *Schriftenreihe des Instituts für Verkehrswesen*, vol. 42, Universität Karlsruhe, 1989.

[23] D. Goldsman, “Simulation output analysis,” *Winter Simulation Conf.*, pp. 97–103, Dec. 1992.

[24] J. M. Charnes, “Statistical analysis of output processes,” *Winter Simulation Conf.*, pp. 41–49, Dec. 1993.

Linda Briesemeister received her graduate degree in Computer Science from the Technical University of Berlin, Germany in March 1998. Since then, she has worked on inter-vehicle communication at DaimlerChrysler Research and Technology as a research scientist. She is currently a Ph.D. candidate in Computer Science at the Technical University of Berlin. Her research interests include distributed algorithms applied on mobile computing and wireless networks as found in inter-vehicle communication. She is a member of ACM and IEEE.

Günter Hommel received his degree Dipl.-Ing. in Electrical Engineering from T.U. Berlin in 1970. He first joined the Department of Electrical Engineering and then the Department of Computer Science at T.U. Berlin where he received his Ph.D. in Computer Science. In 1978 he joined the Nuclear Research Center in Karlsruhe working in the field of real-time systems. From 1980 he was responsible for a research group in Software Engineering at the German National Research Center for Information Technology (GMD) in Bonn. In 1982 he was appointed professor at the Institute for Computer Science of T.U. Munich, where he worked in the fields of real-time programming and robotics. Since 1984 he is professor at the Institute for Technical Computer Science of the T.U. Berlin. He now is the head of the Real-Time Systems and Robotics Group.