A Study on the Feasibility of Mobile Gateways for Vehicular Ad-hoc Networks

Vinod Namboodiri, Manish Agarwal, Lixin Gao
Department of Electrical &
Computer Engineering
Univ. Massachusetts
Amherst, MA
vnambood, magarwal, lgao@ecs.umass.edu

ABSTRACT

Development in Wireless LAN and Cellular technologies has motivated recent efforts to integrate the two. This creates new application scenarios that were not possible before. Vehicles with Wireless LAN radios can use other vehicles with both Wireless LAN and Cellular radios as mobile gateways and connect to the outside world. We aim to study the feasibility of such global connectivity from the road through simulation of the underlying connectivity characteristics for varying traffic and gateway densities. The connectivity results suggest that each vehicle should be able to connect to at least one gateway for a majority of time. The average path lifetimes are found to be good enough many traditional Internet applications like FTP and HTTP. The effectiveness of the AODV wireless ad-hoc routing protocol over this scenario is evaluated and shown to perform well for the densities considered. However, the routes created by AODV can break very frequently due to the dynamic nature of mobility involved. We introduce a couple of prediction based routing protocols to minimize these route breakages and thus improve performance. These protocols take advantage of some deterministic characteristics of the mobility model to better predict route breakages and take preemptive action.

1. INTRODUCTION

With the advent of Cellular 3G networks and Wireless LANs a new form of wireless paradigm has emerged. The ubiquitous connectivity of Cellular networks and the high data rate offered by wireless LANs make a combination of the two attractive. Recent work on combining these two technologies has focussed mainly on the integration aspects [9,12]. In this case, a Wireless LAN is used where an access point is available and a 3G connection otherwise. This scenario assumes each user has the hardware capability and can afford to have both connections. A more immediately feasible approach would be to provide support for intermediate networks where some users having 3G capability share it with others who just have Wireless LAN capability. We visualize an ad hoc network with 3G gateways providing ubiquitous connectivity to the outside world while still maintaining the high data rate wireless LAN connection among themselves. Such intermediate networks would need a different approach to tackle connectivity issues.

As of now, wireless networks have been rapidly consuming territory both on the home front and work place. That is, a user is covered while at home or while at work. An untapped area is the time spent in vehicles. Connectivity while on the road will be an important application area for wireless networks in the coming years. Developing applications and protocols for this setting is a challenge because the mobility is highly dynamic, and very different from what current wireless networks have been subject to through research and deployment.

For vehicles on the road, stationary gateways at different points of the road can be considered for providing ubiquitous connectivity. But the downside of this is the resources required for installations. Another approach is to make some of these vehicles act as gateways for other vehicles. These gateways must use another wireless, longer range technology which can route packets to the nearest point of interest off the road. Such a technology could be Cellular networks. The gateways would have both Wireless LAN and Cellular interfaces. The nodes which do not have Cellular cards, but just Wireless LAN cards, can use the gateway to send packets to points of interest globally as shown in Figure 1. This requires cooperation among users and sufficient gateway density such that a node can at least reach one for a good portion of its time on the road. This is in fact related to the node and gateway density directly. What gateway density is good for a given node density is an open question and one that we intend to shed light on through this study.

We use a discrete first order markov model to characterize the mobility of nodes and provide a simulation based analysis of the underlying connectivity characteristics. We find that paths are available between nodes and gateways for a high percentage of time and these paths on an average last long enough to carry out most of the applications for which the the Dedicated Short Range Communication (DSRC) standard was envisioned for and also some traditional Internet applications like FTP and HTTP. We follow it up with a demonstration of the performance of an ad

VANET'04, October 1, 2004, Philadelphia, Pennsylvania, USA.

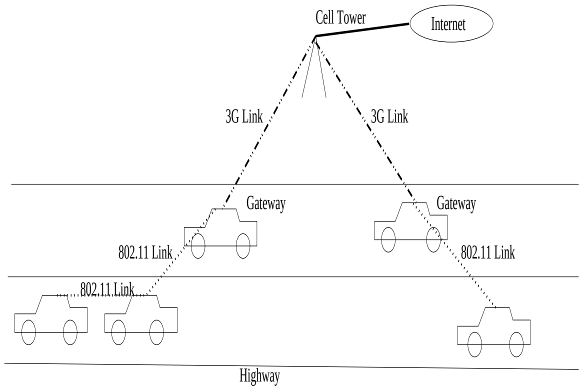

Figure 1: Mobile Gateway Scenario

hoc routing protocol called AODV over the vehicular setting [8]. For the node and gateway densities considered, AODV is able to deliver a high percentage of packets from nodes to their gateways. Though a path may be available, the dynamic mobility pattern causes route breakages very often. It is critical for such dynamic settings that existing paths be fully utilized with minimum packet loss before the network is partitioned. To reduce the ill-effects of frequent route breakages, thus increasing routing performance, we have present two prediction based routing protocols which we have developed. These protocols take preemptive action by predicting when a route would break and try to create an alternate route before the route actually breaks. This results in a better packet delivery ratio.

This paper is divided into the following sections. Section II is a short survey on related work in the area. Section III describes the mobility model we use to characterize the motion of vehicles over a highway. Section IV describes the simulation scenario used for this work. The underlying connectivity of the mobility model is characterized by certain metrics followed by their evaluation in Section V. The performance of AODV for the same scenario is evaluated in Section VI. In Section VII we present our prediction based protocols and evaluate their performance with comparisons to AODV. We finally conclude the paper with Section VIII.

2. RELATED WORK

Ad hoc routing protocols have basically bifurcated into location based protocols and reactive or proactive routing protocols [8,13]. Location based routing is seen as a solution to the issue of scalability, but the traditional reactive routing protocols are still considered better for small scale networks with path lengths of only a few hops. The FleetNet project has done an extensive amount of work on Inter-Vehicular Communication (IVC) with some focussing on connecting vehicles to the internet through stationary roadside gateways [2,6,7,10]. However, they base their work only on location based routing protocols. That is a valid argument when it comes to IVC scenarios where vehicles far apart are trying to communicate with each other. In this paper we are concerned with routing from a vehicle to a gateway vehicle which is expected to be only a few hops away. The highway scenario is expected to be highly partitioned and the probability of forming long paths is small. Thus the issue of scalability does not come into the picture and a reactive routing protocol is expected to do well. Moreover, our objective is to study the feasibility of communication between vehicles and mobile gateways. The IMPORTANT framework compares different ad hoc routing protocols against a suite of mobility models, one of which is a freeway [5]. Its main purpose was to compare the effects of different mobility patterns on routing performance rather than study the freeway model characteristics.

One of our contributions is studying the effectiveness of prediction based routing for such dynamic networks as vehicular ad hoc networks. Previous work on such prediction based schemes had never subjected their protocols to such a mobility model [11,12,15,16]. Our prediction model is an effort to see how such protocols match up to AODV and whether deterministic mobility patterns in vehicle mobility can be exploited to increase their utility. The utility of such models is far greater in vehicular settings where routes break very often and repair and reconstruction could take up valuable time which is at a relative premium compared to other less-dynamic mobility models. Network partitions occur often and fully utilizing the ‘un-partitioned’ time is critical.

Most studies so far have attempted to provide global connectivity to wireless networks through a stationary gateway to the Internet. We believe we are the first ones to study the feasibility of the notion of mobile gateways to ad hoc networks on the road. However, there has been previous work on gateway support to ad hoc networks [7,17,18]. Some of those principles can be used in building mobile gateways.

3. MOBILITY MODEL

Characterizing the motion of vehicles on a road is a difficult task. There can be no one good way to do that. It depends on the layout of the road, the traffic density, and of course the behavior of the drivers. What is important is that some basic properties of the motion of vehicles are captured so that our routing protocols can be tested on those. Vehicular traffic theory basically divides models into macroscopic and microscopic models, with the former concerned with aggregate traffic and the latter based more on individual car mobility which exhibits both spatial and temporal dependence. For this work, we concern ourselves only with the temporal dependence and assume that each node moves independently of others. The traffic densities we consider for studying the effect of gateway densities are relatively sparse where spatial independence is highly probable. Temporal dependency, leading to smooth transitions between speed states, is the most important property which we felt that had to be captured. That is, vehicles should not abruptly move from a very slow speed to a very fast speed. They have to accelerate through several intermediate speeds, in a step by step fashion if discretized. We have tried to capture this pattern in terms of a first order markov chain as shown in Figure 2. Vehicles are modelled as moving from one state (or speed) to another with a certain probability, in terms of p and q in this case. p characterizes the affinity of a driver to a certain speed state and controls how often he stays in a state and moves over to other states. q is derived from p and signifies a divide between lower set of speeds and a higher set of speeds. We discretize speeds into four preferred speeds of 40, 50, 70 and 80 mph (17.8, 22.3, 31.3 and 35.6 m/s). To give each driver a different mobility pattern or signature, an offset v is added to these set of four preferred speeds. This offset, taking both positive and negative values, is dis-

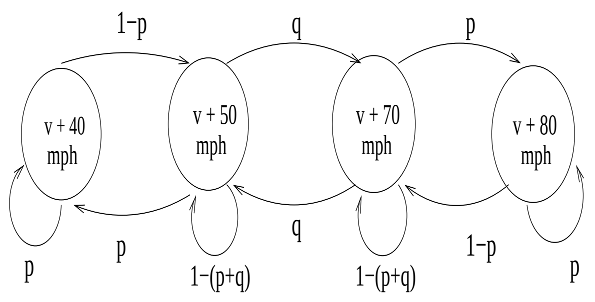

Figure 2: Mobility Model

tributed among the nodes to follow a normal distribution as suggested by classical vehicular traffic theory [14]. We use a distribution with a mean of about 60 mph and a 85th percentile value of about 70 mph. These values can be corroborated from the following sources [3,4]. The value of p and, hence q, is found from the stationary distributions of the above markov model. Solving the transition matrix of the model yields,

$$\pi_1, \pi_4 = p/2$$

and

$$\pi_2, \pi_3 = (1-p)/2$$

where

$$\pi_i,\ i = 1, 2, 3, 4$$

is the stationary distribution of each state of the markov model. π_1 and π_4 represent the probabilities of states $v+40$ and $v+80$ while π_2 and π_3 are the probabilities for states $v+50$ and $v+70$ of the stationary distribution. This stationary distribution is also the long run distribution of the model.

The vehicles (referred to also as nodes throughout this paper) calculate a 'next speed' and move at that speed for the next INTERVAL before recalculating a new next speed. The INTERVAL value can be increased or decreased to change the rate at which nodes could possibly change speeds. To study if there is any effect on communication between vehicles with varying v values, we simulate three variants of the model. The offset v is related to p and by varying range of v, we let p values vary from 0.25 to 0.35, 0.2 to 0.4 and 0.15 to 0.45. This simulates different degrees of speed variability among vehicles. We term these variants as 2535, 2040 and 1545 respectively.

4. SIMULATION SETTING

The highway scenario we consider is a long straight stretch of road. For our work, we are concerned with a 20 km segment only. Nodes are uniformly distributed along the segment. Nodes moving out of the segment are replaced by new nodes at the beginning of the segment. These new nodes are given the same state as the nodes that exited the 20 km segment. This helps in maintaining the average density throughout the simulation. We also try to simulate route disconnections due to nodes moving out through exits on the highway. All nodes are given exponentially distributed lifetimes, when initialized, with an average equal to the length of the simulation run. This value was chosen so that only a moderate amount of nodes exit during the simulation. At the end of their lifetimes, these nodes move out of the road at their minimum preferred speeds. Again to keep things stable, these nodes are brought back in from a different point at their lowest preferred speeds. We do not study the effects of exiting directly. The performance results we show later factor in the effect of nodes exiting.

Experiments are carried out with INTERVAL set at 10 seconds. We base this study only on one direction of traffic. We consider traffic densities of 1 per 100m, 1 per 150m and 1 per 200m. We vary the density of gateway nodes from 1 per 400m to 1 per 700m in steps of 100m. We will refer to all these densities as 100, 150, 200, 400, 500, 600 and 700. The gateway nodes are mobile nodes with all functionalities of the normal nodes. In our work, we primarily use these as destinations for packets of traffic flows which could possibly be routed further through to the cellular network.

Five nodes are designated as source nodes and send packets to any one of the gateways with preference given to gateways that are closer to the source node. We do not study the issues of handoffs here.We use a ON-OFF traffic model with spurts of 20 seconds followed by a pause of 30 seconds before the next ON spurt. 64 byte CBR packets are sent at a rate of four packets per second. We use a ON-OFF model as it is more representative of the type of traffic likely in the scenario we consider. The small number of sources was also intentionally chosen to focus more on routing and less on load effects. The NS2 simulator was used for the evaluation of routing protocols [1].

5. CONNECTIVITY METRICS AND EVALUATION

Characterizing the connectivity of vehicles is very important in understanding what is expected from any routing protocol. The notion of connectivity at the routing layer depends mainly on reachability at the physical layer. A path exists between two nodes only if there is no partition in the network between them. When considering a road in a single direction, this boils down to whether the length of the separation edge (link joining two clusters) in the corresponding network graph exceeds R, the range of radio. We use the following metrics to describe the connectivity of vehicles with special emphasis on connectivity to vehicles that act as gateways. Vehicles which are not gateways are termed as nodes.

5.1 Metrics

The following metrics are those that we believe to be most important for the concept of mobile gateways on the road. A thorough understanding of these should give an idea of the feasibility of mobile gateways for providing global connectivity on the road.

5.1.1 Path Lifetime

This metric shows the average time a unbroken path existed between two nodes. Simply stating, it is the average time the nodes spent within a network with partitions on either side. As long as no partition separates them, there is logically a path between them. We use this metric to show how long, on an average, paths existed between nodes and gateways. The average is taken only for paths which existed at least once between a node and a gateway. It can be argued that path breakage may not affect applications if a new one can be found soon enough. However, session switching may not be cost effective. Also, with heavy load, packets will be still dropped aplenty when buffers overflow.

5.1.2 Gateway Connectivity Ratio

This metric shows the average percent of time where at least one gateway was reachable. This is a most important metric as being connected to multiple gateways at one time and connected to none at other times will severely hamper performance. Ideally, each node is able to reach at least one gateway for as much percent of time as possible. The average is taken over the length of simulation period.

5.1.3 Number of Gateways

This metric points out the number of gateways to which a node had a path to during the simulation period. This basically shows the number of choices a node has to choose from when selecting a gateway. The effect on this metric by varying node density and gateway density should be interesting.

5.1.4 Gateway Lifetime

Though the path lifetime metric showed how long a unbroken path lasted between a node and a gateway, it did not show how long a node had a path to that gateway. Gateways can be disconnected and reconnected. This metric takes reconnections also into account and shows how long a node had a path to a gateway even if it was broken many times.

5.2 Evaluation

The statistics are collected for source and gateways for each variant of p values. These are then averaged across node and gateway densities. Standard deviations are shown for each mean value and these project the variability of values among different source-gateway pairs.

5.2.1 Path Lifetime

Figure 3 shows how the path lifetime varies over different node and gateway densities. The important thing to note is that most values were in the range of 30-40 seconds. This means that on an average, a route can be constructed between a node and gateway for about 35 seconds before they become disconnected due to a partition. For the kind of applications DSRC is envisioned for, this amount of path lifetime seems reasonable. A few examples are emergency vehicle warnings, traffic congestion reporting and some safety applications. Traditional Internet applications like FTP, HTTP e.t.c which do not require continuous connectivity for long periods of time at a stretch should also be successful. From Figures 3 and 4 it can also be noticed that the path lifetime seems to actually increase with decrease in node and gateway density. The reason for this is that the values shown are averages for different node-gateway pairs and as mentioned earlier, only paths which ever existed are used for calculations. As the node density decreases, only a few paths are formed and these are between the nodes that are in close proximity which leads to longer path lifetimes. At a high node density, more paths are formed,with some between distant nodes, but as a result incur more frequent breakages. As expected, the main difference between the variants is that 1545 seems to show more variability in its values compared to the other two.

5.2.2 Gateway Lifetime

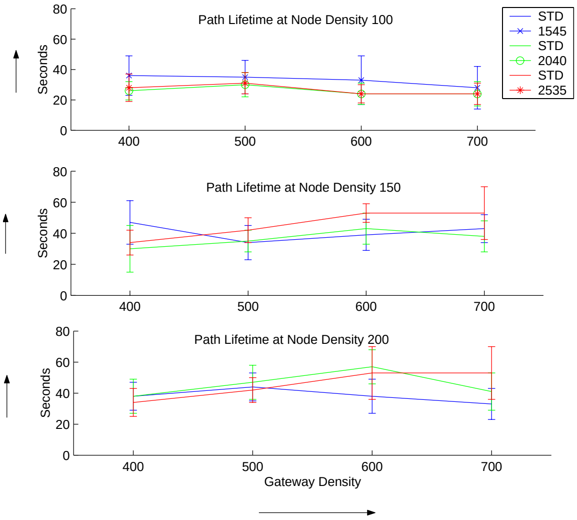

Figure 3: Path lifetime for different node and gateway densities

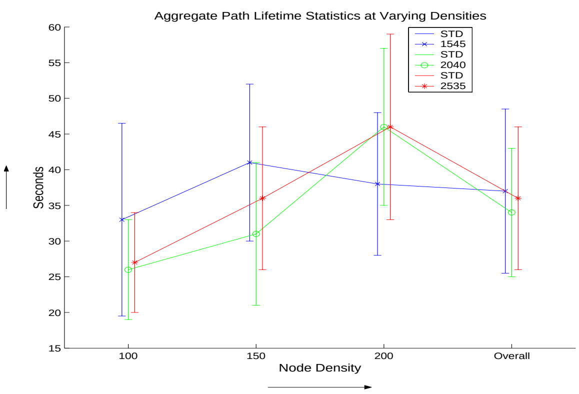

Figure 4: Path lifetime aggregate statistics for different node densities

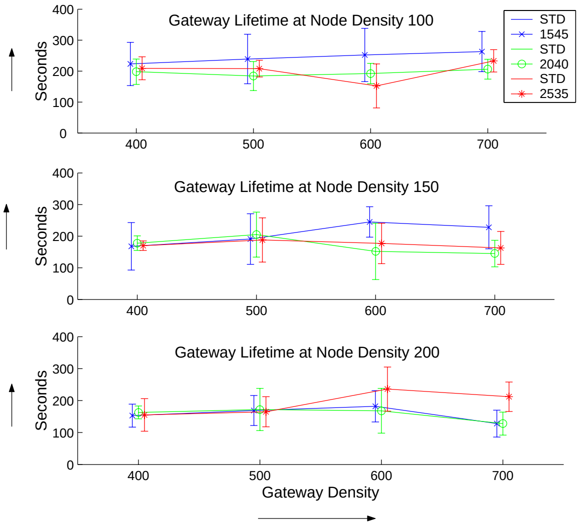

Figure 5: Gateway lifetimes for different node and gateway densities

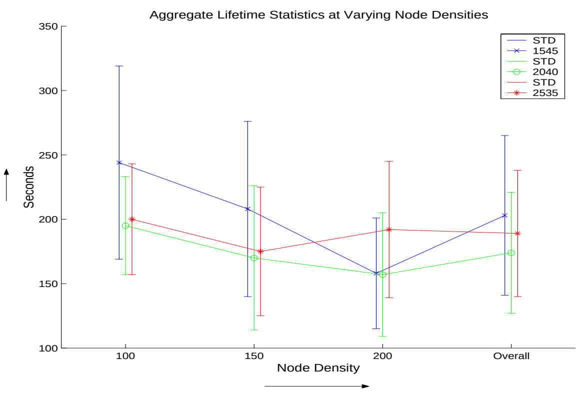

Figure 6: Gateway lifetimes aggregate statistics for different node densities

The Gateway lifetimes, as shown in Figures 5 and 6, decrease with node density as lesser paths are maintained. However, the statistic that a node maintains a path with a gateway for about 200 seconds, even if the path breakage time is much lesser, can have significance for certain applications where some state has to be stored for a certain time after which it can be passed to another node or sent over to the wired network through the cellular network.

5.2.3 *Number of Gateways*

As expected, the number of gateways available to a source decreases with decrease in gateway density as shown in Figures 7 and 8. This phenomenon repeats for decreasing node density too because of lesser number of paths towards gateways. From the perspective of handoff protocols, dealing with how to choose a gateway will be important, given that at high node and gateway densities, there are usually multiple options. This could also be important from a load balancing point of view.

5.2.4 *Gateway Connectivity Ratio*

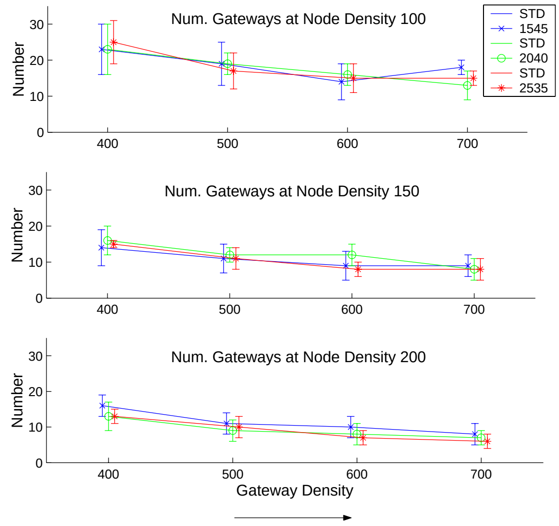

Figure 7: Number of gateways for different node and gateway densities

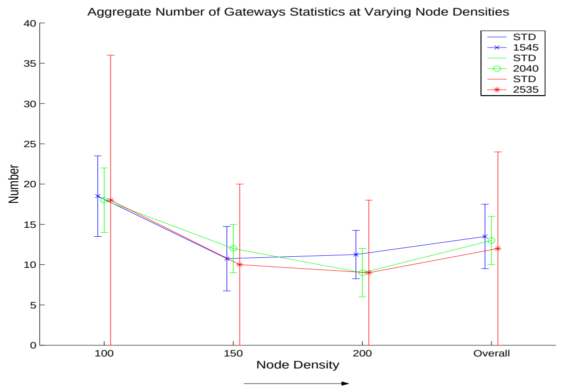

Figure 8: Number of gateways' aggregate statistics for different node densities

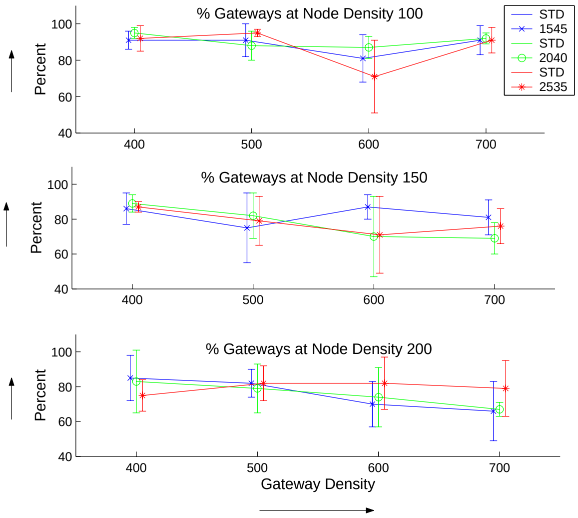

Figure 9: Percentage of time gateways were available for different node and gateway densities

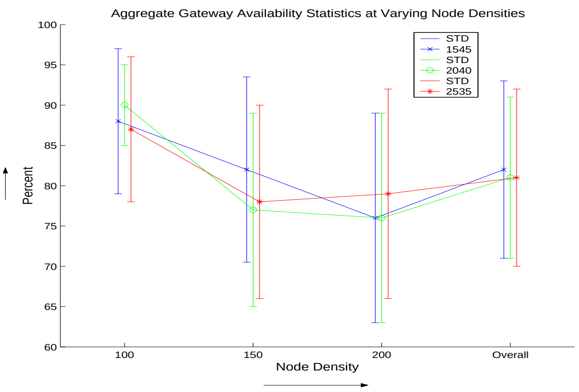

Figure 10: Gateway availability aggregate statistics for different node densities

Figure 9 shows that there is some decrease in gateway connectivity with decrease in gateway density. There seems to be an even greater impact due to decrease in node density (Figure 10). On an average there is about 80-85% gateway connectivity across different node densities which is pretty good considering we are focusing on only densities representing free-flowing independent traffic and have not considered congested scenarios where vehicles could be packed quite closely.

6. ROUTING PERFORMANCE EVALUATION USING AODV

Evaluating Ad Hoc On-Demand Distance Vector Routing (AODV) for supporting mobile gateways would be a step towards understanding the requirements to support such mobile gateways. There have been efforts in this direction for stationary gateways before [17].

AODV is a reactive routing protocol and creates routes on-demand only. A route is created uses a resource request (RREQ) and resource reply (RREP) sequence between source and destination or intermediate nodes which have a route to the destination. If the source gets back multiple routes, it chooses the one with minimum hop count.

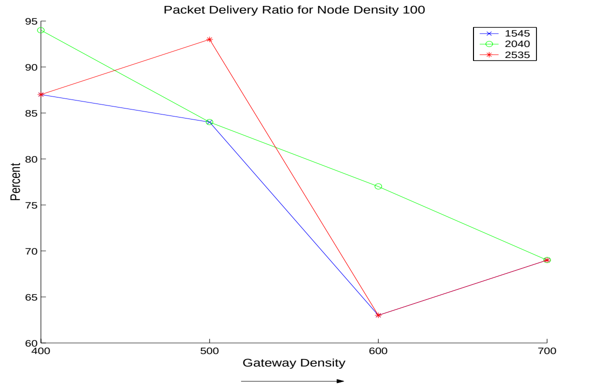

Figure 11: Packet delivery ratio statistics for a node density of 100

6.1 Packet Delivery Ratio

The main metric we show for AODV is the packet delivery ratio which is the ratio of data packets sent to data packets received over all source gateway connections. As shown in Figures 11, 12, 13, the packet delivery ratio tends to decrease significantly with decrease in gateway density. Figure 14 shows the packet delivery ratio aggregated for each node density irrespective of gateway densities. There seems to be some variability in this trend with decrease in node density. This is due to the fact that in some scenarios even when densities are small, some sources are fortunate to have a gateway and hence are able to have a good routing performance. Also the traffic model plays a part. If gateways are unreachable more often during the OFF spurts, this will not be reflected on the routing performance results. At high node densities the routing performance seems pretty reasonable given the dynamic scenario. Some loss recovery mechanisms at the application layer and retransmissions could pretty much ensure all packets are delivered for these high node densities. For lower node densities more drastic measures may need to be taken like a node multicasting its packets to multiple gateways and hoping at least one of them is able to successfully receive an enough percentage of its packets. Other schemes may also be employed, a discussion of which is outside the scope of this paper. When one looks at the aggregate packet delivery statistics it is also apparent that the different variants of p do not seem to have much difference. This is possible because, as explained above, a lot depends on instantaneous connectivity characteristics when packets are being sent. More work is required to understand exactly how variability between nodes affects routing performance.

6.2 Path Lengths

The average hop counts involved for all scenarios varied between 1 to 2 hops. In other words, for the node and gateway densities considered, the routing protocol usually found the nearest gateway within an average of 2 hops. A small hop count didn't always mean a good packet delivery ratio however. Some of the small hop counts was because longer paths couldn't be formed and hence no packets could be delivered. As Figure 15 shows, the average hop counts tended

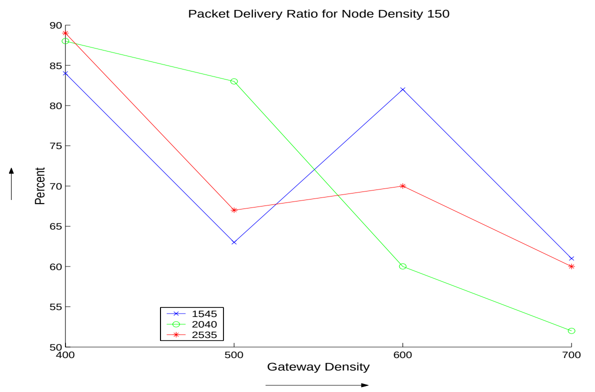

Figure 12: Packet delivery ratio statistics for a node density of 150

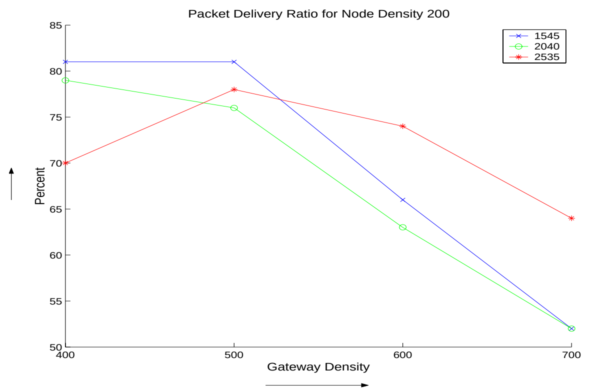

Figure 13: Packet delivery ratio statistics for a node density of 200

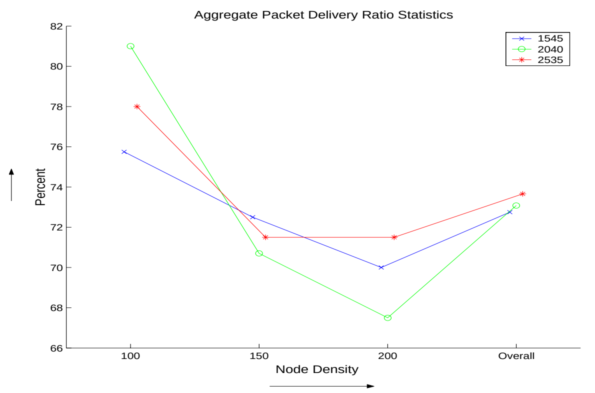

Figure 14: Packet delivery aggregate statistics for different node densities

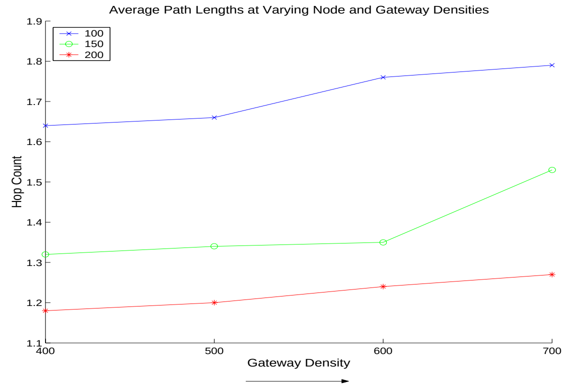

Figure 15: Average path lengths for varying gateway densities

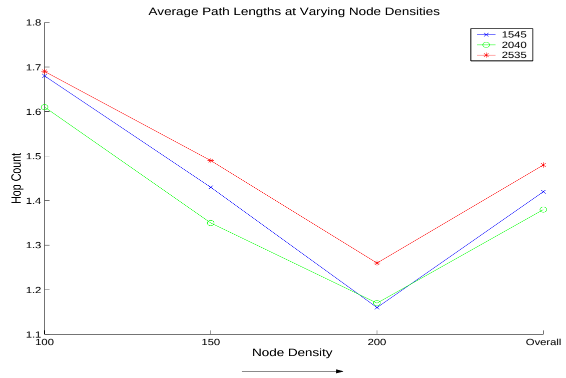

Figure 16: Average path length for varying node densities

to increase with increase in gateway density as for fewer gateways, longer paths were required. However, as node density decreased (Figure 16), the longer paths were never found leading to a decrease in average path lengths. The small path lengths involved justify our decision to use a reactive routing protocol rather than a location based routing protocol.

6.3 Route Lifetimes

This metric shows how often routes break in the dynamic mobility scenario. To calculate this metric we could not use the traffic model used so far as it used ON-OFF spurts which would have masked some of the route breakages. Hence, we just simulated AODV over the same three variants 1545, 2040 and 2535 with a traffic model which sent CBR packets continuously for 400 seconds. The notion of gateways was also not used as it would have led to switching of paths pretty often. We designated five nodes as sources and sent packets to a common destination. We found that the routes created by AODV broke about every 5 seconds on an average. A connected path existed between these source-destination pairs for much longer but the links along the route broke and a new route had to be constructed every time. This motivates the idea of using route prediction to minimize the ill-effects of frequent route breakages by preemptively creating new routes before the old ones break.

7. PREDICTION BASED ROUTING

We have developed two prediction based protocols called PRAODV and PRAODV-M which use the notion of link and route lifetime estimates. These are variants of AODV designed to take advantage of some deterministic characteristics of the motion of vehicles on the road. Compared to mobility models like the random waypoint model [8], motion of vehicles have a deterministic direction and certain maximum bounds on acceleration. Vehicles do not abruptly change speeds. They move through speeds in a smooth fashion. These facts can be used to predict how long a link between two vehicles will last. Prediction based protocols are likely to be most useful in situations where route breakages are frequent. We use distance as the primary indicator of connectivity. In the real world, a more complex function is required for indicating connectivity. We present our work using only distance because of its simplicity and use it to show how useful prediction based routing can be. In the real world, more complex functions of distance or SINR or both would be required.

7.1 PRAODV and PRAODV-M Protocols

PRAODV retains most of the features of the AODV protocol. The main modification is in the RREP reply packet sent from the destination or intermediate nodes to the source. Whenever a node sends a reply, it includes its velocity and location information in the packet. Every subsequent node that receives this reply on route to the source of the request makes a link lifetime prediction based on its own location and velocity and the values inside the reply packet from the node that sent it. It adds its predicted link value to the the reply packet replacing the old predicted value if its estimation of the lifetime of link is lesser than any previous estimations of any link of the route. It also replaces the location and velocity information of the previous node with its own values before forwarding it towards the source. The basic idea is to have an estimation value which is the minimum of all links along the route. This is the predicted lifetime of the route. A new request is sent out just before the end of this predicted lifetime to construct a new route to the destination. The idea is not to wait till packets are dropped due to route failure and try to construct a new alternate route where possible. This is the only difference between AODV and PRAODV, with both using minimum hop count as the metric to choose between multiple paths to the same destination.

The other variation of prediction based AODV routing, PRAODV-M, uses the path which has the maximum predicted value among multiple route options as metric unlike AODV and PRAODV which use minimum hop count. The idea behind this is to minimize preemptive route creation by choosing the route which is expected to last the longest. How well the lifetime of a route can be estimated plays a key role in the performance of this protocol.

7.2 Prediction Methodology

A route between any two nodes consists of $n-1$ links with n the number of nodes forming the route including the source and destination. A route breaks when any one of these $n-1$ links break. Hence a route is only as strong as its weakest link. A link breaks when the two nodes on either side of it move out of the communication range, R (links could also break due to obstructions but that is outside the

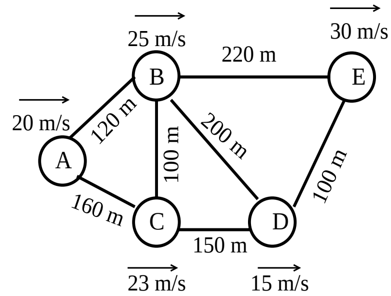

Figure 17: Sample scenario for using prediction

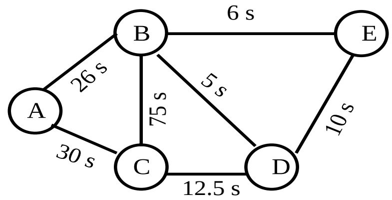

Figure 18: Corresponding predicted lifetime values

scope of this paper). Hence if two nodes are at a distance d_{ij} from each other, $R - |d_{ij}|$ represents the absolute distance the nodes have to separate additionally in order for the link to break. Thus if the two nodes have velocities V_i and V_j then the absolute difference in velocities is represented by $|V_{ij}|$. Thus the lifetime of a link can be predicted as

$$Pr.\ Link\ Lifetime_{ij} = \frac{R - |d_{ij}|}{|V_i - V_j|}, \quad V_i \neq V_j \qquad (1)$$

Hence,

$$Pr.\ Route\ Lifetime_{sd} = Minimum\ Pr.\ Lifetime_{ij},$$

among all links from source s to destination d.

Two vehicles with very similar speeds can inflate the predicted value of that link. Hence we cap this predicted lifetime with a maximum value to avoid such unrealistically high values.

Figures 17 and 18 show a sample scenario and corresponding predicted values calculated based on equation (1) to give an idea of the way prediction is done. If A seeks a route to E, AODV will choose AE as this is the route with shortest hop count. PRAODV uses the same route AE and predicts the lifetime of the route to be 6 seconds as it is the bottleneck lifetime value among the two links AB and AE. It sends a route request just before 6 seconds have elapsed, even though the route may still be up. PRAODV-M looks at all the options available to it and finds that path ACDE has the highest lifetime value of 10 seconds among all paths from A to E and uses that route to send packets from A to E. It sends a new route request just before 10 seconds have elapsed to look for a new route.

7.3 Evaluation

For comparing the three routing protocols, we use a traffic model where 5 source nodes send 64 byte CBR packets to a common destination node for 400 seconds at the rate of four packets per second. The hop lengths of the paths used by all three protocols for all variants varied between 1 and 3 hops.

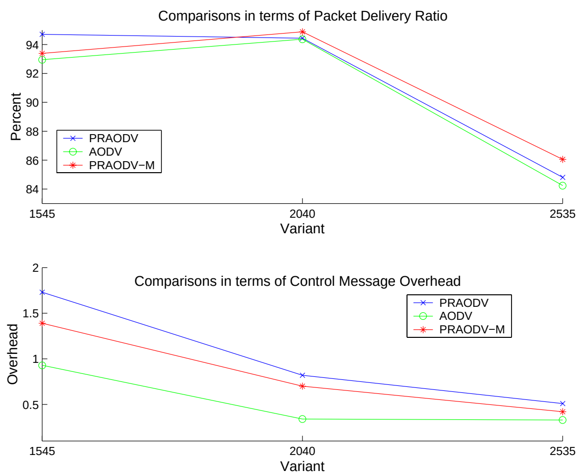

Figure 19: Comparisons in terms of packet delivery ratio and overhead

As Figure 19 shows, there is improvement in the packet delivery ratio on using PRAODV and PRAODV-M over AODV. The margin of improvement of PRAODV over AODV is based purely on the fact that in AODV, each time a route breaks, a new one has to be constructed before packets can be sent through. This leads to some packets being dropped. PRAODV will always do equal or better than AODV at the cost of overhead as discussed below. The margin of improvement of PRAODV-M over AODV depends on how accurate the predicted lifetimes are. For PRAODV-M to do better than AODV, the routes chosen using predicted route lifetimes as metric should be better than using hop count as metric. Since the prediction is done only when the route is constructed, speed changes in nodes could in fact lead to routes chosen by AODV to be better. So there are cases where AODV performs slightly better than PRAODV-M. But we found that PRAODV-M performs better in a majority of the cases.

As expected AODV was the best in terms of overhead as shown in Figure 19. The prediction based protocols caused more control messages per packet delivered by preemptively creating routes. However, the usage of predicted lifetimes as metric in PRAODV-M has decreased the overhead compared to PRAODV. PRAODV-M selects the route with highest predicted lifetime and waits till just before this lifetime expires before trying to create a new route. This leads to fewer control messages being sent out.

More accurate prediction will further reduce the overhead incurred. One approach to do this would be to predict lifetimes when actual data packets are being sent and not rely entirely on the predictions made during route construction. We intend to pursue this and similar approaches in future work for vehicular settings. In fact, overhead may not be that big an issue for vehicular ad hoc networks. They are not power-constrained like many wireless devices are. The bandwidth provided by an 802.11 connection far exceeds that of a cellular link. Thus the bottleneck will always be the uplink from the mobile gateway to the Cellular basestation with plenty of spare capacity expected on the 802.11 links.

8. CONCLUSIONS

This paper has extensively studied the connectivity between nodes and mobile gateways for vehicular ad-hoc networks through simulation. A simple, but quite representative mobility model of vehicles on the road was presented and was used as the basis for the study. Various aspects of connectivity along with routing protocol performance were evaluated. What can be expected for different node densities and gateway densities was described. The connectivity metrics indicated that the path lifetimes are long enough for many applications foreseen by the DSRC standard and other traditional Internet applications. For reasonable node and gateway densities, a vehicle was found to be able to connect to a gateway for a good percentage of the time it is on the road. The performance of the AODV routing protocol also indicates that the idea of mobile gateways has a good chance to succeed in providing global connectivity to vehicles on the road. To offset the frequent route breakages on using AODV, we present and evaluate two prediction based variants of AODV. These take advantage of deterministic mobility patterns of vehicles on the road. The improvement in packet delivery ratio using these protocols was weighed against the cost of overhead. With overhead not being as major a concern in vehicular networks, these protocols could have great utility.

Acknowledgment

This work has been supported by the NSF Grant CCR-0329794. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

9. REFERENCES

[1] "The network simulator ns-2. http://www.isi.edu/nsnam/ns2,".
[2] "The FleetNet Project," http://www.et2.tu-harburg.de/fleetnet/english/documents.html
[3] "Speed Trend Characteristics," http://www.ibiblio.org/rdu/gov-stat/speed-tr.html
[4] "Traffic Speeds on English Trunk Roads:2003," Transport Statistics, Department for Transport, UK.
[5] F. Bai, N. Sadagopan, and A. Helmy, "The IMPORTANT Framework for Analyzing the Impact of Mobility on Performance of Routing for Ad Hoc Networks," *AdHoc Networks Journal - Elsevier Science*, Fall 2003.
[6] M.Bai, L.Wolf, and W.Franz, "Mobile Internet Access in FleetNet," *13. Fachtagung Kommunikation in Verteilten Systemen KiVS 2003*, Leipzig, Germany, February 2003.
[7] M. Bechler, O. Storz, W. Franz, L. Wolf, "Efficient Discovery of Internet Gateways in Vehicular Communication Systems," *57th IEEE Semiannual Vehicular Technology Conference VTC 2003-Spring*, Jeju, South Korea, April 2003.
[8] J. Broch, D.A. Maltz, D.B. Johnson, Y-C Hu, and J. Jetcheva, "A Performance Comparison of multi-hop wireless ad hoc network routing protocols," In *Proceedings of MOBICOM 1998*, pages 85-97, October 1998.

[9] M. Buddhikot, G. Chandranmenon, S. Han, Y. W. Lee, S. Miller, and L. Salarelli, "Integration of 802.11 and Third-Generation Wireless Data Networks," In *Proceedings of IEEE INFOCOM 2003*, April 2003

[10] H. Füessler, M. Mauve, H. Hartenstein, M. Ksemann, and D. Vollmer, "A comparison of routing strategies for vehicular ad-hoc networks," In *Proceedings of MOBICOM 2002*, Student Poster, 2002.

[11] H.-Y. Hsieh and R. Sivakumar, "Internetworking WWANs and WLANs in Next Generation Wireless Data Networks," *International Conference on 3G Wireless and Beyon*, San Francisco, CA, USA, May 2002.

[12] T. Goff, N. B. Abu-Ghazaleh, D. S. Phatak and R. Kahvecioglu, "Preemptive Maintenance Routing in Ad Hoc Networks," *MOBICOM 2001* , Rome, Italy, July 2001.

[13] M. Mauve, J. Widmer, H. Hartenstein, "A Survey on Position-Based Routing in Mobile Ad-Hoc Networks," In *IEEE Network. Vol. 15 (2001)*, Nr. 6, S. 30-39.

[14] M. Rudack, M. Meincke and M. Lott, "On the Dynamics of Ad-Hoc Networks for Inter-Vehicle Communications (IVC)," *The 2002 International Conference on Wireless Networks ICWN 2002*, Las Vegas, USA, June 2002

[15] S. Su and M. Gerla, "Wireless ad hoc multicast routing with mobility prediction," *Mobile Networks and Applications*, vol. 6, no. 4, pages 351-360, 2001.

[16] W. Su, S.-J. Lee, and M. Gerla, "Mobility Prediction and Routing in Ad Hoc Wireless Networks," *International Journal of Network Management*, 2000.

[17] Y. Sun, E.M. Belding-Royer, and C.E. Perkins, "Internet Connectivity for Ad hoc Mobile Networks," *International Journal of Wireless Information Networks special issue on "Mobile Ad Hoc Networks (MANETs): Standards, Research, Applications*", 9(2), April 2002

[18] R. Wakikawa, K. University, J.T. Malinen, C.E. Perkins, A. Nilsson, and A.J. Tuominen, "Global connectivity for IPv6 mobile ad hoc networks," *IETF Internet draft*, draft-wakikawa-manet-globalv6-01.txt (2002).