Prediction Based Routing for Vehicular Ad Hoc Networks*

Vinod Namboodiri, Lixin Gao
Department of Electrical and Computer Engineering
University of Massachusetts at Amherst
Amherst, MA, U.S.A
Email: {vnambood, lgao}@ecs.umass.edu

Abstract

Development in short range wireless LAN (WLAN) and long range wireless WAN (WWAN) technologies have motivated recent efforts to integrate the two. This creates new application scenarios that were not possible before. Vehicles with only WLAN radios can use other vehicles that have both WLAN and WWAN radios as mobile gateways and connect to the Internet while on the road. The most difficult challenge in the scenario is to deal with frequent route breakages due to dynamic mobility of vehicles on the road. Existing routing protocols which are widely used for mobile ad hoc networks are reactive in nature and wait till existing routes break before constructing new routes. The frequent route failures result in significant amount of time needed for repairing existing routes or reconstructing new routes. In spite of the dynamic mobility, the motion of vehicles on highways is quite predictable compared to other mobility patterns for wireless ad hoc networks, with location and velocity information readily available. This can be exploited to predict how long a route will last between a vehicle requiring Internet connectivity and the gateway which provides a route to the Internet. Successful prediction of route lifetimes can significantly reduce the number of route failures. In this paper we introduce a prediction based routing (PBR) protocol that is specifically tailored to the mobile gateway scenario and takes advantage of the predictable mobility pattern of vehicles on highways. The protocol uses predicted route lifetimes to pre-emptively create new routes before existing ones fail. We study the performance of this protocol through simulation and demonstrate significant reductions in route failures compared to protocols that do not use pre-emptive routing. Moreover, we find that the overhead of pre-emptive routing is kept in check due to the ability of PBR to predict route lifetimes.

---

*A preliminary version of this work appeared in *Proceedings of Vehicular Ad hoc Networks VANET'04* [1].

I. INTRODUCTION

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 and potentially a selling point for cars which have the capability to do so. Such a technology will enable gaming, multimedia streaming and other Internet applications to be carried out while traveling in the vehicle proving to be a great value added service. It could also provide access to a great deal more safety applications than currently available and may in fact even cut some production costs by reducing the number of safety features required on board. Developing applications and protocols for this setting is a challenge because the mobility is highly dynamic, and very different from what most current wireless networks have been subject to through research and deployment.

Any approach to achieve connectivity while on the road requires considerations of bandwidth, cost, and seamless mobility. Currently some efforts are underway to provide cars with broadband access coinciding with the emergence of new wireless Wide Area Network (WWAN) technologies like 3G, 4G, WiMax and satellite based services. These technologies work over large ranges (order of miles or more) and can provide seamless connectivity at moderately high bandwidths. Unfortunately, it may not be affordable for every vehicle to be enabled with such technology. It may also be overkill in some situations when all that is desired is short range communication for some safety applications. IEEE 802.11 or Wireless LAN (WLAN) technology is a mature, high-bandwidth and low cost technology with the capability to organize into ad hoc multi-hop networks. Due to its low-cost, it is feasible for every vehicle to have an 802.11 network interface and has already been adopted for some short range safety applications. However, due to their limited range (order of hundreds of meters only), WLANs cannot by themselves enable ubiquitous connectivity to vehicles.

Traditionally, the concept of wireless connectivity from a vehicle has been looked at from the perspective of inter-vehicular communications (IVC) [2]. It is only recently that the research community has focused on the task of providing Internet connectivity to vehicles [3], [4]. Their approach has been to deploy static gateways alongside roads where vehicles connect to these gateways wirelessly (802.11) to gain access to the global Internet. The static or installed gateway approach entails a great deal of investment which may only drive up the cost of the eventual broadband access available rendering it as expensive as other longer range technologies. Moreover, as vehicles move from one area to another, gateways will have to be switched every few seconds driving up the overhead to maintain routes to

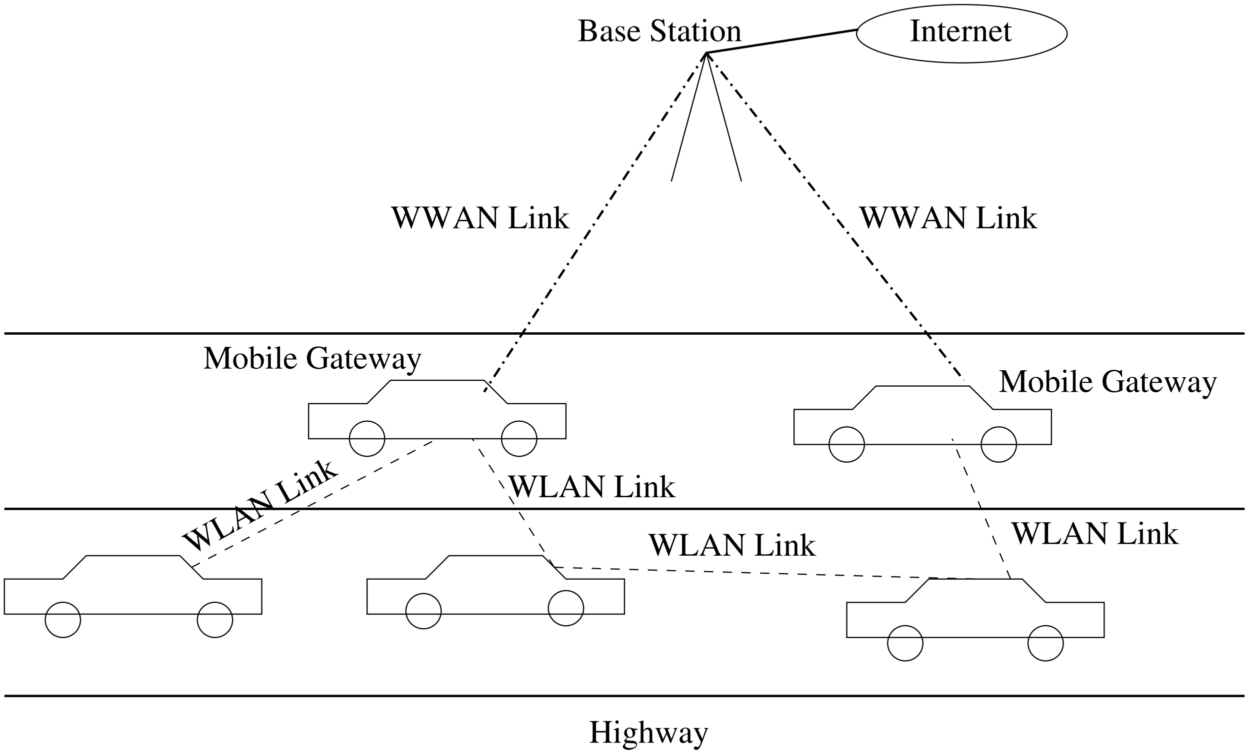

Fig. 1. Mobile Gateway Architecture

gateways. The use of access points already deployed in urban communities is an alternative way to provide connectivity from vehicles [5], [6]. This approach, however, has the same issues of very frequent gateway switching as described above, with the possibility of lack of coverage in many areas without usable wireless access points.

The mobile gateway approach is based on the theory that inexpensive IEEE 802.11 WLANs can supplement longer range technologies to bring down the cost for Internet access for many vehicles and its comparatively high bandwidth could prove useful for short range applications. It would be an incrementally deployable solution where some vehicles with both WLAN and WWAN communication capabilities act as mobile gateways for other ‘node’ vehicles with just a WLAN interface, providing ubiquitous connectivity to the outside world while still maintaining the high data rate WLAN connection among themselves. The 802.11 WLAN multi-hop networking capabilities can be fully leveraged to provide seamless connectivity across different gateway and node densities. Mobile gateways do not require any infrastructure based support like static gateways do, and can provide access without being restricted by any geopolitical boundaries. Also, the fact that gateway vehicles are moving in the same direction reduces the relative motion between node vehicles and their gateways. This provides for much fewer gateway switches and resulting in smaller overhead than a static gateway approach.

With the advantages of such a scenario fairly obvious, the next question then is its feasibility given the current state of the art in wireless networking and vehicle densities on the roads. Such a combination of technologies has already been researched for traditional wireless network scenarios [7], [8]. However, a lot of work needs to be done to adapt them for the highly dynamic mobility of vehicles on highways, where routes break frequently degrading application performance. Nevertheless, there are advantages that can be derived from the somewhat predictable motion of vehicles on the road to develop routing protocols for the scenario. Vehicles move along pre-defined paths, i.e. roads, and their speeds are restricted to within certain bounds. This provides an opportunity to predict how long routes would last compared to arbitrary motion patterns like the random waypoint model [9].

This paper presents a prediction based routing protocol, PBR, that predicts how long routes will last and pre-emptively creates new routes to replace old ones before they break. Our main contributions are the following: (i) A mobility model that characterizes individual as well as collective motion of vehicles on a free-flowing highway. This model abstracts out the important details of highway mobility patterns enabling the study of various factors important to routing protocols for the scenario. (ii) The PBR protocol specifically tailored to the highway vehicular mobility pattern to reduce the frequency of route failures. Through our performance evaluation we find that the PBR protocol is able to take good advantage of the relatively predictable nature of highway mobility patterns to reduce route failures significantly without using too much control overhead. This overhead is kept in check due to the ability of PBR to use predicted route lifetimes to pre-emptively create new routes before existing routes fail. (iii) The study of vehicle densities and required fraction of gateways to achieve sufficient connectivity to meet application requirements. We find that freeways within urban areas in the U.S have sufficient vehicle density for the mobile gateway scenario to be feasible. Moreover, at these densities the required fraction of vehicles that need to be gateways is less than 20%.

The remainder of this paper is organized as follows. Section II describes routing protocols for wireless ad hoc networks and how our work relates to existing work. In Section III we present our highway mobility model. In Section IV we introduce our PBR protocol along with the mechanisms used to obtain location and velocity information and how this information is used to predict route lifetimes. In Sections V and VI we describe our performance evaluation setting, simulation experiments and their corresponding results. Concluding remarks are made in Section VII.

II. RELATED WORK

Traditionally mobile ad hoc routing protocols have been classified as proactive, reactive or location based [9]–[11].

In a proactive routing protocol, all nodes send routing messages at pre-determined periods to create new routes or update existing ones [12], [13]. Proactive protocols suffer from the fact that it is difficult to determine what messaging period is best to maximize routing performance while avoiding excessive control overhead. These protocols are also classified as table-driven protocols since they try to maintain updated routing tables.

On the other hand, reactive protocols create routes on an on-demand basis with no new routes sought until existing ones break [14]–[16]. This keeps the control overhead low, but lacks sensitivity towards new, better routes which might be available with time; frequently the case for vehicular mobility patterns. Thus, for dynamic mobility patterns, the number of route failures is high, requiring repair or reconstruction often. However for the mobility models considered generally for mobile ad hoc networks, like the random waypoint model, the low control overhead of the reactive protocols is found better suited for power constrained devices like laptops and PDAs.

Location based protocols which route using only location information have been proposed for IVC [2], [3], [17]. These protocols rely on a location server to provide location information of all destination nodes. Maintaining such location servers for vehicles on the road is difficult, incurring a lot of overhead to maintain it with current information. Moreover, the advantage with regard to scalability that these protocols have may not be important in the mobile gateway scenario on highways, where routes to gateways, if present, are expected to be of a few hops in length only.

We propose the PBR routing protocol specifically tailored to the highway mobility scenario that improves upon routing capabilities of a reactive protocol without using the overhead of a proactive protocol. The deterministic motion pattern and speeds of vehicles can be exploited to predict roughly how long an existing route between a ‘node’ vehicle and a ‘gateway’ vehicle will last. By creating a new route just before the predicted route lifetime expires, the idea is to preempt route failures and make the most of the connectivity available leading to smaller network downtime. By making such ‘educated’ guesses about when nodes should try to create new routes, the control overhead is expected to be much smaller than that of proactive protocols accompanied by a lower network downtime compared to reactive protocols which need to repair and reconstruct broken routes. The motivation is that in the mobile gateway scenario, Internet connectivity may not always be available due to absence of routes to any gateways. It is thus imperative to utilize the available time as efficiently as possible.

Approaches to predict route lifetimes and use them for preemptive routing have been used before in the general context of mobile ad hoc networks, but never for vehicular ad hoc networks [18], [19]. The mobile gateway scenario presents new challenges which require an in depth study of when a prediction based protocol is effective and when it is not based on the mobility pattern as well as node and gateway density required for connectivity.

Most studies so far have attempted to provide global connectivity to wireless networks through a stationary gateway to the Internet [2], [20], [21]. Some of those principles can be used in building mobile gateways for the scenario considered by our work.

III. HIGHWAY MOBILITY MODEL

Here, we describe our mobility model for the highway scenario. It is important that the model is able to abstract out the important details of vehicular mobility that are important for ad hoc routing.

A. Highway Mobility Model

The motion on highways is characterized by high speeds with all vehicles within certain maximum and minimum speed bounds. The maximum bound is enforced by the speed limit for highways. The minimum bound corresponds to a rough estimate of how slowly some vehicles would move on a highway or some minimum speed limit [22]. Thus, most vehicles are moving with similar speeds most of the time with occasional acceleration or deceleration. This results in certain periods where links of a route to a gateway remain connected after which existing links break requiring new routes to be created.

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 flow theory and simulation models are typically classified as macroscopic or microscopic [23], [24]. When following a macroscopic approach, one focuses on system parameters like traffic density (number of vehicles per mile 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 on a stretch of road. In general, from a macroscopic perspective, vehicular traffic is viewed as a fluid and existing fluid models are applicable. In contrast, with a microscopic approach, the movement of each individual vehicle is characterized with spatial and temporal dependence being important characteristics. In order to generate vehicle movement patterns for wireless ad hoc routing experiments, one clearly has to follow a microscopic approach since the position of each individual vehicle is needed to determine whether a pair of vehicles can communicate with a certain range of radio communication.

Our highway mobility model is a discrete time model with a fast lane and a slow lane with vehicles recalculating their acceleration every Δt seconds. For each vehicle i, it's next speed depends on its current speed by the relation

$$v_i(t + \Delta t) = v_i(t) + a_i(t).\Delta t\,. \tag{1}$$

Vehicles are assigned to the slow lane or fast lane randomly. The acceleration for each vehicle i at time t, $a_i(t)$ is calculated using the equations below. U_1, U_2, U_3 and U_4 are uniformly distributed random variables between 0 and 1. Each vehicle has a range of acceleration, A_{max}, and de-acceleration, D_{max}, from which it randomly picks a value every certain step of time or it decides to not accelerate at all.

$$a_i(t) = \begin{cases} U_2\ .A_{max}, & \text{if } U_1 < acc_i + p_r \\ U_2\ .(-1).\ D_{max}, & \text{if } acc_i + p_r \leq U_1 < acc_i + dacc_i + 2p_r \\ 0, & \text{otherwise,} \end{cases} \tag{2}$$

where $acc_i + p_r$ is the probability of a vehicle i to accelerate and $dacc_i + p_r$ is the probability of a vehicle i to decelerate. acc_i and $dacc_i$ are probabilities for a driver of vehicle i to accelerate or decelerate based on individual behavior. These are assigned as

$$acc_i \ = \begin{cases} U_4\,(1 - 2p_r), & \text{if } U_3 < 3AGG/4 \\ 0, & \text{otherwise,} \end{cases} \tag{3}$$

$$dacc_i \ = \begin{cases} U_4\,(1 - 2p_r), & \text{if } 3AGG/4 \leq U_3 < AGG \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

Parameter AGG controls the aggressive behavior of drivers by giving them higher probabilities to accelerate or decelerate by specifying the fraction who have such behavior[1]. Highway traffic studies suggest that about 75% of aggressive drivers tend to favor acceleration over a general mean velocity, which is used above for assigning values of acc_i and $dacc_i$ [22]. Additionally, all vehicles randomly accelerate or decelerate with a probability p_r each. This provides a random motion to vehicles that can represent momentary speed changes and spatial dependencies between vehicles. We do not directly correlate the speed of a vehicle to that of its surrounding vehicles. The assumption of two lanes and a

[1] By aggressive we mean tendency of drivers to prefer either acceleration or deceleration compared to just a general random motion around the mean velocity

free-flowing highway scenario coupled with the indirect use of p_r to capture spatial dependence helps validate the individual spatial independence within the model. The higher the value of p_r, the higher is the inclination of vehicles to change speeds.

The motion of each vehicle i has a behavioral component acc_i or $dacc_i$ controlled by parameter AGG as well as a random component p_r which is also a parameter of the model. Thus, the model can be configured to adjust the number of aggressive drivers by assigning suitable values (between 0 and 1) to AGG with the randomness of the mobility pattern controlled by p_r.

Additionally,

$$\text{if } v_i(t) > V_{max}, \;\; v_i(t) = V_{max}, \tag{5}$$

or

$$\text{if } v_i(t) < V_{min}, \;\; v_i(t) = V_{min}. \tag{6}$$

There are maximum and minimum velocity bounds, V_{max} and V_{min}, which restrict the velocity within that range. Moreover, every vehicle that has a velocity greater than $(V_{max} + V_{min})/2$ moves on the fast lane while a vehicle below this level moves on the slow lane.

We believe that such a mobility model is quite representative of the highway scenario with similar models used in previous wireless ad hoc network studies [25], [26]. The speed limits and driver behavior statistics used by our model have been adopted from [22]. Focusing only on free flowing highway traffic allows us to ignore intersections and crossings, which would be better modeled by a traffic simulation package [27], [28]. This simplifies the model without sacrificing the important details of the highway mobility pattern. We are mainly concerned with the general state of highway motion ignoring outliers like exiting or entering the highway, impact of construction work and similar events. The ability to control each driver's motion pattern as well as the general motion of all vehicles enables the study of highway mobility patterns. Our model gives control over how predictable the mobility pattern is by varying p_r over a reasonable range.

IV. PREDICTION BASED ROUTING PROTOCOL (PBR)

Two important building blocks of our PBR protocol are obtaining location and velocity information of vehicles on the route to the gateway, and the prediction algorithm that uses this information to predict when the route will break. An additional consideration is whether to use routes through vehicles that are not moving in the same direction as the vehicle requiring the route, the source node. Current routing protocols for wireless ad hoc networks do not generally differentiate among nodes based on what direction they are moving. This is because, either the mobility pattern to which these protocols are subjected to have no specific direction or the protocol has no use for such information. The direction in which nodes move is very important for the lifetime of a link between two nodes. A link formed between two nodes moving in opposite directions is expected to be of a much shorter duration than if they were moving in the same direction. In the highway mobility pattern, nodes move either in the same direction or opposite direction (ignoring curved roads for simplicity). This information about direction can be easily used in selecting gateways and forwarding nodes for a route, and predicting route lifetimes. If connectivity to gateways is sufficient using just vehicles moving in the same direction as the source node, then routes through oncoming vehicles should be avoided. However, in cases where the vehicles density is not sufficient or there are not enough gateways moving in the same direction, routes through oncoming vehicles can be used. We start with describing the basic operation of the protocol.

A. Basic Operation

The PBR protocol's basic operation of creating routes is similar to that of a reactive protocol. When a node needs to communicate to a location on the Internet it checks its routing table for a route. If none exists, it broadcasts a route request (RREQ) packet with a time to live (TTL) value specifying the number of hops to search for a gateway that would have the required route. The RREQ packet contains a sequence number, source ID, destination ID, source node's direction of motion, and a list of nodes and their direction of motion appended by the intermediate nodes as the message travels through the network. A neighbor upon receiving this packet forwards it to all its neighbors after reducing the current TTL value by one if a) the packet contains a higher sequence number for the same source and destination pair and TTL value in the packet is greater than one or if b) the packet contains the same sequence number as a previously received packet for the same source destination pair, but this time with all intermediate nodes on the route traveling in same direction and with the TTL value in the packet greater than one. Condition b) is required so that preference is given to routes having all intermediate nodes traveling in the same direction compared to those routes which may have one or more intermediate nodes moving in the opposite direction. Once such a packet is received by a node, it drops all further packets unless they have a higher sequence number. Imposing the above conditions on packet forwarding helps in mitigating the broadcast storm problem in wireless ad hoc networks [29].

When the RREQ packet reaches a gateway with the desired route to the sought destination, a route reply (RREP) packet is sent back to the source node by the gateway using the chain of nodes in the RREQ packet[2]. This route through which the RREP packet traversed is stored as the source node's gateway route. The source node then uses this supplied route for sending its application packets. If multiple gateways reply, the source node uses the following order of preference: (i) choose the gateway which is at a shortest distance from it (in terms of hops), and has all nodes on the route (including the gateway) moving in the same direction as itself (ii) choose the gateway which is nearest in terms of hops. If the source gets multiple routes for the same gateway, it uses the route which has the maximum predicted route lifetime. Intermediate nodes do not reply back to the source even if they have a route to a gateway, as motion is very dynamic and stale routes increase route failures. The gateway sends back packets from the Internet to the source node using the best route from among all the RREQ packets it got from that node with the highest sequence number (using same criteria as above to select among multiple routes to a source node). In case that route fails, it discovers a new route to the source as above with the source node as the destination of the RREQ.

The PBR protocol also differs from reactive protocols like AODV or DSR in the way it proactively creates new routes before they break. The RREP packet from a gateway to a source in conjunction with the prediction algorithm (explained below) is used to give the source a predicted lifetime for the route. A timer is started at the source equal to this lifetime, with the source sending out a new RREQ for a gateway just before this timer expires. This new RREQ creates a new route (if possible) replacing the old route (if different). Thus, every source node tries to pre-empt route failure by using predicted route lifetimes to proactively seek possibly better routes. This approach works much better than proactively creating new routes after specified intervals because the pre-emption interval is now adaptive based on how long each route is expected to last.

Pre-emptive route creation based on predicted route lifetimes continues until no route can be found since the routing protocol has no idea when the application layer will be sending packets. Maintaining active routes in this fashion, however, wastes control overhead if there is no data to send. To overcome this, PBR protocol adapts based on when the last packet was sent by the source node. If the last packet was sent before a certain time threshold `pred-timeout`, pre-emptive route creation is turned off. The value of `pred-timeout` can be set to make the protocol behave like a reactive protocol in scenarios where data traffic is rare. Whenever a source has a packet to send and finds the route broken, either due to over-estimation of route lifetime or for any other reason, a new route is created with a RREQ-RREP exchange as described above with the pre-emptive timer restarted.

---

[2] The source routing technique for wireless ad hoc networks was first introduced by the DSR protocol [14].

B. Obtaining Route Lifetime

After a RREQ packet is sent, the RREP packet from the gateway is used to gather information to predict the lifetime of the route. Initially the gateway node puts its location and velocity information and sets a *lifetime* field in the RREP packet header equal to some value which is expected to be greater than the minimum of all link lifetimes along the route; a parameter `maxlifetime`. As the packet traverses back to the source, each intermediate node does the following. Based on the velocity and location information of its predecessor (available in RREP packet) and that of itself, it predicts the lifetime of the link between the two nodes using the prediction algorithm. If this value is smaller than the current *lifetime* mentioned in the RREP packet, the *lifetime* field of the packet is changed to this value. Else, the current value in the packet is left unchanged and forwarded towards the source. In this manner, when the source gets an RREP back, the value of the *lifetime* field indicated in the packet is the route's predicted lifetime. Location and velocity information is assumed to be available to each node, which is quite feasible for modern day vehicles equipped with GPS or other navigational instruments.

C. Prediction Algorithm

Let the range of communication of the Wireless LAN technology used be R. If the absolute distance between two nodes i and j is denoted by $|d_{ij}|$, and their corresponding velocities given by v_i and v_j respectively, the lifetime of a link between i and j is predicted as,

$$Lifetime_{link} = \frac{R - |d_{ij}|}{|v_i - v_j|}. \tag{7}$$

Since a route comprises of one or more links, the route lifetime is the minimum of all its link lifetimes. i.e.

$$Lifetime_{route} = \min_{\forall links \in route} \{Lifetime_{link}\}. \tag{8}$$

This minimum value is obtained directly from the RREP packet as explained before. Figures 2 and 3 show sample calculations of an arbitrary scenario. For example, the lifetime for the route A-C-D-E is 10 seconds. The parameter `maxlifetime` is used to place an upper bound on the predicted lifetime of a link. This is necessary for cases where the velocities between two nodes are very close to each other, which could result in the value of predicted link lifetime to blow up to large values. Note that in Equation 7 we ignore the width of the road for simplicity. This results in only a slight under-estimation of link lifetimes.

Equation 7 does not take into account the fact that one of the vehicles might be moving closer towards the other. In such a case, the predicted link lifetime should be greater. Figure 4 shows an example. Equation

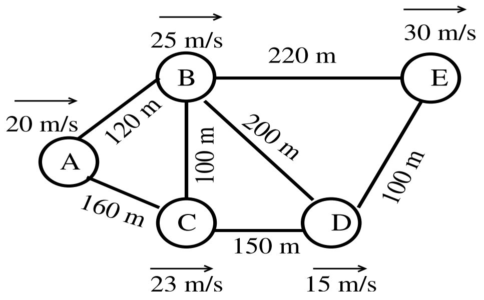

Fig. 2. Example scenario

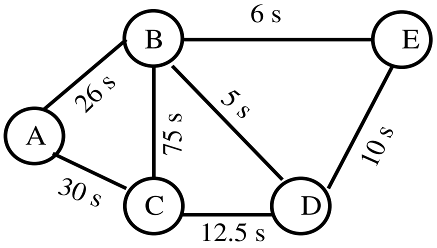

Fig. 3. Calculation of predicted lifetime with R = 250 m

7 would have predicted link lifetime between A and B to be $(200-140)/6 = 10$ seconds. However, since B is behind A and has a speed greater than A, the link lifetime can be predicted higher. For such cases, our algorithm adds an extra `large-bonus` seconds to the lifetime of a link if the difference in speeds between the nodes is greater than `speed-diff` or `small-bonus` seconds otherwise. The value of `speed-diff` is used in the decision because a greater speed difference indicates a greater likelihood that the vehicle approaching from behind will overtake the vehicle ahead. Note that, these bonus times are used in this scenario, instead of modifying Equation 7, because we find that a conservative lifetime prediction is more beneficial in cases where link lifetimes can be high.

The link lifetime equation is altered slightly in the algorithm to predict link lifetimes when routes pass through vehicles moving in opposite directions as well. Since the road carrying oncoming traffic can be quite far apart from the road carrying forward traffic, the algorithm now takes into account this separation distance, w. As before, we ignore the width of the road in this calculation for simplicity. Thus, if two end vehicles of a link are moving in opposite directions,

$$Lifetime_{link} = \frac{\sqrt{R^2 - w^2} + s.\sqrt{d_{ij}^2 - w^2}}{v_i + v_j}, \tag{9}$$

where $s = 1$ when the two vehicles are moving towards each other, and $s = -1$ when they are moving

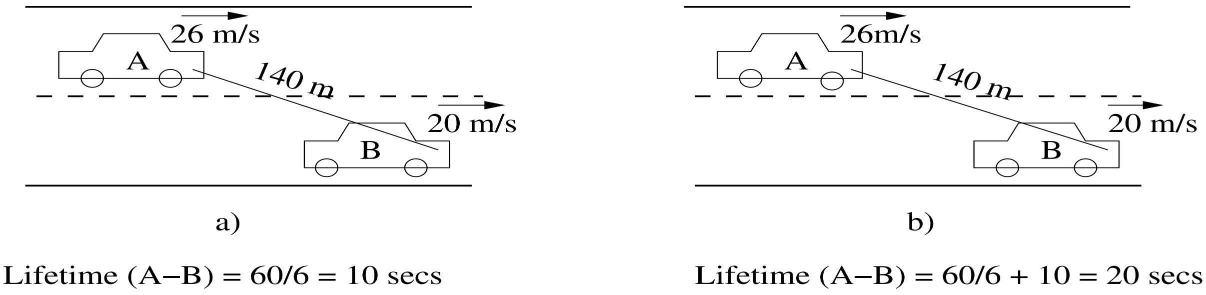

Fig. 4. Prediction taking into account relative positions and speeds, R = 200m a) Prediction that disregards the fact that nodes are moving closer together b) Prediction that adds extra lifetime when nodes are moving closer

away from each other. This lifetime gives the expected time based on current speeds for two vehicles at a distance d_{ij} from each other before they move apart by more than the communication range R. Note that $d_{ij} \geq w$ always for a constant separation distance between a pair of straight roads.

Distance between two nodes is used as the primary indicator of whether two nodes are able to communicate with each other. When two nodes move apart by a distance greater than R, their link is assumed to be broken. This model can be easily adapted to provide more sophisticated predictions based on radio characteristics, signal strength, direction of motion among others, and applied to other wireless ad hoc network settings as well. Similar to using routes through oncoming vehicles, vehicles moving in other directions or other static Wireless LAN installations can be considered too when faced with a lack of vehicle density resulting in inadequate connectivity. The protocol and prediction algorithm can be modified as above for these scenarios too.

D. PBR Variants

The PBR protocol as described above finds the nearest gateway (based on order of preference described above) it can, and routes through it. Thus, when an old route breaks, a new route is constructed to the nearest gateway regardless of whether a route could be constructed to the previous gateway. This approach assumes that gateway switching can be handled seamlessly. However, switching gateways may not always be easy to handle at the application layer depending on the type of application. To take this into consideration we modify PBR such that its goal is to stick with a gateway as much as possible to avoid gateway switching. In this variant, termed PBR-S, once a gateway is selected for a node, the same gateway is sought by the node until no route can be constructed to it. Then onwards, a new gateway is sought as in the PBR protocol.

Another variant, PBR-M, selects the gateway among all those within a certain number of hops with the largest predicted route lifetime. Once a route breaks, PBR-M selects a new gateway with the same largest predicted lifetime criteria. Thus, this variant makes use of the ability to predict route lifetimes in selecting gateways to which routes will exist longer without failure. This is expected to result in longer lasting routes between the source node and gateway, and also reduce the number of gateways switched as a result.

We explore the effectiveness of these variants compared to the PBR protocol in our performance evaluation.

V. PERFORMANCE EVALUATION SETTING

The primary aim of the performance evaluation of PBR protocol is to identify the effect of density of nodes and gateways and mobility patterns on routing performance through simulation. We also look at how the control overhead of using prediction based pre-emptive routing compares to that of proactive and reactive protocols. For that we first describe our primary routing metrics and the simulation environment followed by the mobility model and PBR protocol parameters we employ.

A. Routing Metrics

We use the following as our main routing metrics. The average of the metric for all the nodes was used to form the metric value for a single experiment.

1) Packet Delivery Ratio: The ratio of data packets delivered at the destination to those generated by the CBR sources.

2) Route Failures: The number of times existing routes failed and new ones had to be created. This metric shows how effective a protocol is in avoiding route failure. In the evaluation section we specifically look at the percentage of dropped packets due to route failures rather than specifically the number of route failures. This is so as to abstract out the effect of route failures removing the dependence on the length of the simulation.

3) Route Requests Sent: The PBR protocol tries to minimize route failures by pre-emptively creating new routes. This metric gives an account of the control overhead incurred by the routing protocol by counting all the attempts made at creating routes, successful or otherwise.

B. Simulation Environment

To evaluate the performance of the PBR protocol, we constructed a packet-level simulator which allowed us to observe and measure the protocol's performance under various conditions. The experimental setting used a 2000m long straight stretch of highway with two lanes for vehicles to move. Only one direction of motion was considered for the experiments we present. Where applicable, we point out the results of considering both directions of motion. The stretch of road was simulated as a torus so that any vehicle exiting the area was brought back in from the beginning. The torus scenario was important to avoid the border effect for nodes at the ends. The border effect creates isolated nodes at the ends due to lack of other nodes outside the border. All nodes and gateways are randomly placed on the road with each assigned one of the two lanes also randomly[3]. Based on the lane they are placed, all nodes and gateways are given an initial velocity. The initial velocity assigned to each node was the average speed of the lane they were placed in. All nodes re-calculated new acceleration values for themselves every STEP seconds. The simulation run time was 1 hour.

The data traffic model we use is CBR (constant bit rate) with a rate of 1 packet per second. Each source node sends packets to its nearest gateway until the route fails. After failure, the source node again finds the nearest gateway (which might be different than the one before failure) and continues. Our packet level simulator does not employ a MAC protocol. It is assumed that when a node has a packet to send, it is able to send it without any interference from other node. The results based on packet delivery ratio we show in our experiments thus need to be adjusted to account for a few extra packet drops due to contention at the MAC layer.

C. Mobility Model and PBR Protocol Parameter Settings

Table I lists the fixed mobility model parameters used in the simulation with Table II listing the PBR protocol parameters. The mobility model parameters chosen are reasonable values based on realistic speed trends [22], [30]. The value of p_r chosen is such that the expected time for non-aggressive nodes to change their speeds is 10 seconds based on $STEP = 5$ seconds, which makes for a highly random mobility pattern that can test the PBR protocol well. We further study the effect of varying p_r in a separate experiment in the evaluation section. The PBR parameters are chosen based on what we deemed reasonable based on the experimental mobility and traffic setting. Different settings might require different suitable values. The wireless communication range R was set to 200m based on IEEE 802.11 radio characteristics for outdoor environments.

---

[3] We also study a non-random scenario in Section VI-D.

TABLE I

MOBILITY MODEL PARAMETERS

Maximum Velocity	$70\,mph = 31.3\,m/s$
Minimum Velocity	$40\,mph = 17.8\,m/s$
Maximum Acceleration (A_{max})	$5\,m/s^2$
Maximum Deceleration (D_{max})	$5\,m/s^2$
$STEP$	$5\,s$
Δt	$1\,s$
AGG	0.2
p_r	0.25
Highway Length	$2000\,m$
Range, R	$200\,m$

TABLE II

PBR PROTOCOL PARAMETERS

MaxLifetime	$50\,s$
Pred-Timeout	$25\,s$
Small Bonus	$2\,s$
Large Bonus	$10\,s$
Speed-Diff	$5\,m/s$
Max. # RREQ Retries	3

VI. EXPERIMENTS AND RESULTS

A. *Effect of Vehicle and Gateway Density on Connectivity*

Our first endeavor is to study vehicle and gateway densities for the scenario that provides adequate connectivity[4]. We define connectivity as the percentage of time at least one gateway is reachable from a node. Depending on the location of the highway stretch which could be in an urban or rural area, the total vehicle density would vary. Thus, under some circumstances, there may not be enough gateways to provide adequate levels of connectivity to the Internet. Thus, we would like to know, given a certain vehicle density on a highway, how many of these vehicles have to be gateways to achieve the desired levels of connectivity. Without an adequate vehicle and gateway density, no routing protocol can deliver

---

[4] The number of vehicles is the sum of nodes and gateways. Thus by looking at only vehicle and gateway densities, we can infer node density characteristics as well.

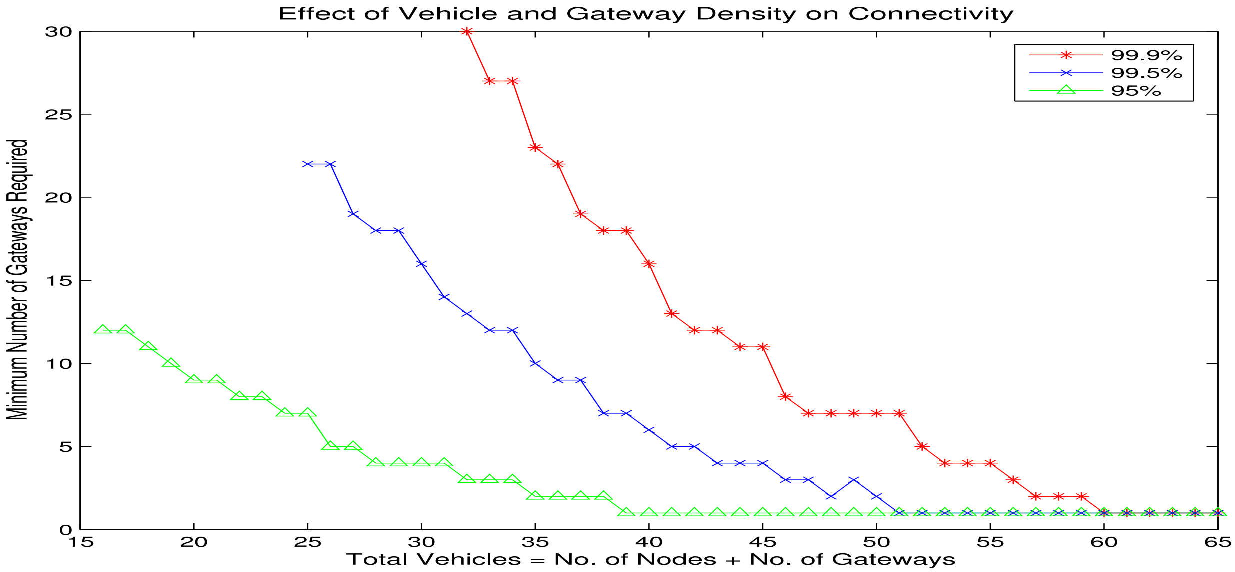

Fig. 5. Effect of vehicle density on minimum number of gateways required to achieve a certain desired level of connectivity

the desired application performance. We consider three levels of connectivity in this experiment, namely 99.9%, 99.5% and 95%. For some application scenarios a gateway must almost always be reachable, while for some others small periods of lack of connectivity may be acceptable. The three levels we consider should cater to these varying desired levels of connectivity. The plot shown in Figure 5 is obtained by a breadth first search technique on the induced graph of the underlying topology for every second of the simulation time. This enabled us to find the average number of seconds for all nodes when a gateway was reachable or in other words the connectivity as a percentage of total simulation time.

Based on the plot, it seems that the number of gateways required decreases exponentially as the vehicle density increases. If the vehicle density is greater than 60 vehicles per 2000m (which corresponds roughly to a vehicle every 33m), then theoretically the minimum number of gateways required is only 1 even for the highest levels of connectivity desired. Thus, for ‘high’ vehicle densities, the fraction of vehicles that are required to be gateways is very small. Larger number of nodes require only a small number of gateways due to the multi-hop capabilities of wireless ad hoc networks.

For low vehicle densities, a large fraction of vehicles should be gateways. For ‘very low’ vehicle densities, it does not matter how much fraction of nodes are gateways; the desired connectivity levels will be hard to achieve. For high levels of connectivity (99.9%), the minimum vehicle density required is about 32 vehicles per 2000m which translates roughly to a vehicle every 62m. For low levels of connectivity (95%) the minimum vehicle density comes down to 16 per 2000m which is half of that required for 99.9% connectivity.

Based on the plot, we can conclude that a vehicle every 40m (50 vehicles in our simulation on a 2000m highway stretch) will provide adequate connectivity with only a small fraction of gateways (less than 20%) required. At lower densities, a higher fraction of gateways will be important to satisfy connectivity requirements.

To verify that these vehicle densities are reasonable values to expect on roadways we studied U.S highway traffic statistics [31]. It was found that the vehicle densities on freeways of almost all urbanized areas were above 30 vehicles per 2000m. Moreover, the vehicle densities on freeways of the most populated urbanized areas easily exceeded 60 vehicles per 2000m. For these areas, less than 5% of vehicles of gateways suffice to guarantee 99.9% connectivity. For state run highways and other roads the vehicle densities were 6-10 times smaller. Looking at the statistics based on time of day [32], it was found that the freeways on urbanized areas had sufficient vehicle density from about 6 AM to 10 PM. Thus, the density values from our experiment are reasonable for high vehicle density roads like freeways in an urban area for most of the day time.

With low vehicle density in the forward direction, using routes through oncoming vehicles has an effect similar to doubling the vehicle density in the forward direction in terms of connectivity. But this results in a significantly higher frequency (at least 50% more) of route failures justifying the avoidance of such routes if possible.

B. *Effect of Node and Gateway Density on Routing Performance*

Here, we study how the route length or path length between a node and its gateway varies by varying the number of nodes and gateways and what is the effect on routing performance. As route lengths increase, the chances of route failure should also increase. A longer route with more intermediate nodes is more susceptible to route failures since it requires only one of the links between these nodes to break for the route to fail, the probability of which increases as the number of links on a route increase. Here we look at the number of packets dropped due to route failures for the PBR protocol and characterize the effect of route length. This is important to determine the effect of a certain node and gateway density on routing performance (not connectivity).

For studying effect of node density, we vary the number of nodes from 35 to 55 in steps of 5 with the number of gateways fixed at 10. For studying effect of gateway density, we vary the number of gateways from 10 to 2 in steps of 2 with the number of nodes fixed at 50. These densities are chosen such that

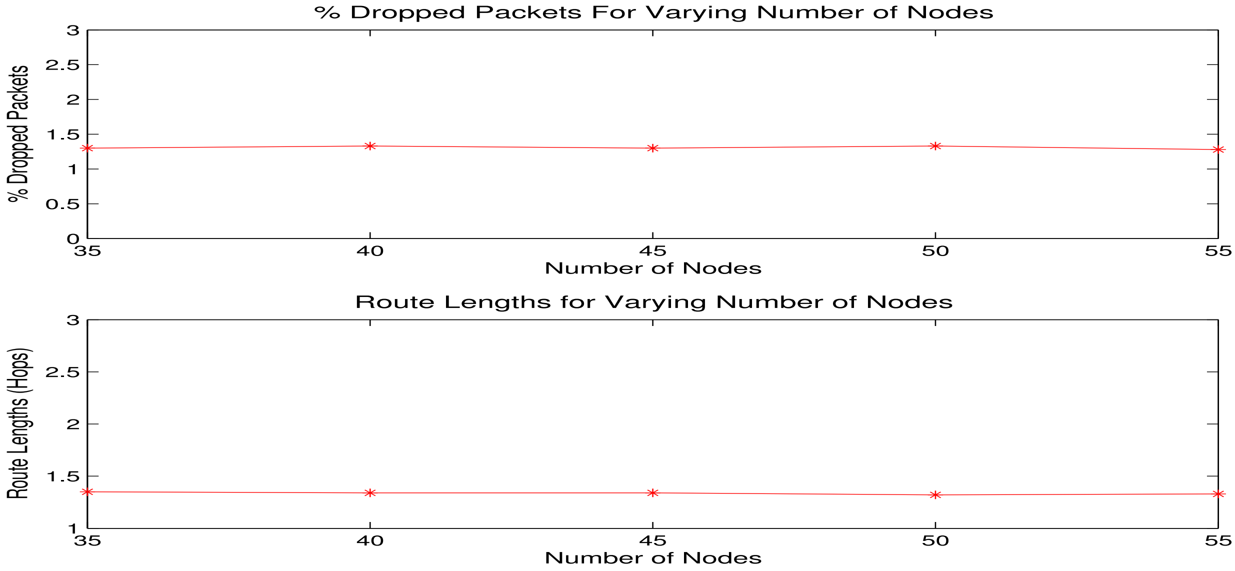

Fig. 6. Effect of varying number of nodes (keeping number of gateways fixed) on route lengths and, thus, dropped packets due to route failures

connectivity for the scenario is very high (refer to previous experiment) allowing us to focus on the impact of route length on route failures.

As Figure 6 shows, node density does not have an effect on route lengths at all and, thus, does not have any effect on route failures. The number of packets dropped due to route failures stays constant and so does the route lengths. On the other hand, Figure 7 shows that gateway density is critical to route lengths and thus route failures. As gateway density decreases, the route length and number of route failures rise.

Larger node density helps in improving connectivity but has no effect on reducing route failures since it does not decrease route lengths. The impact of gateway density on route lengths warrants that the number of gateways be as high as possible to reduce ill-effects of route failures. Routes with gateways at more than two hops from a node suffer a lot of dropped packets due to route failures, thus justifying our policy of choosing the nearest gateway available.

Since smaller route lengths minimize frequency of route failures, we also studied both directions of motion for this experiment, since, there was now a chance to decrease route lengths by also using gateways moving in other direction. Interestingly, the number of packets dropped due to route failures had indeed decreased, but this was because the PBR protocol was more successful in its prediction. The reason for this being that when routes through oncoming vehicles were utilized, the predicted route lifetimes were much smaller than before, increasing the chances that the vehicles will keep their vehicle speeds the same

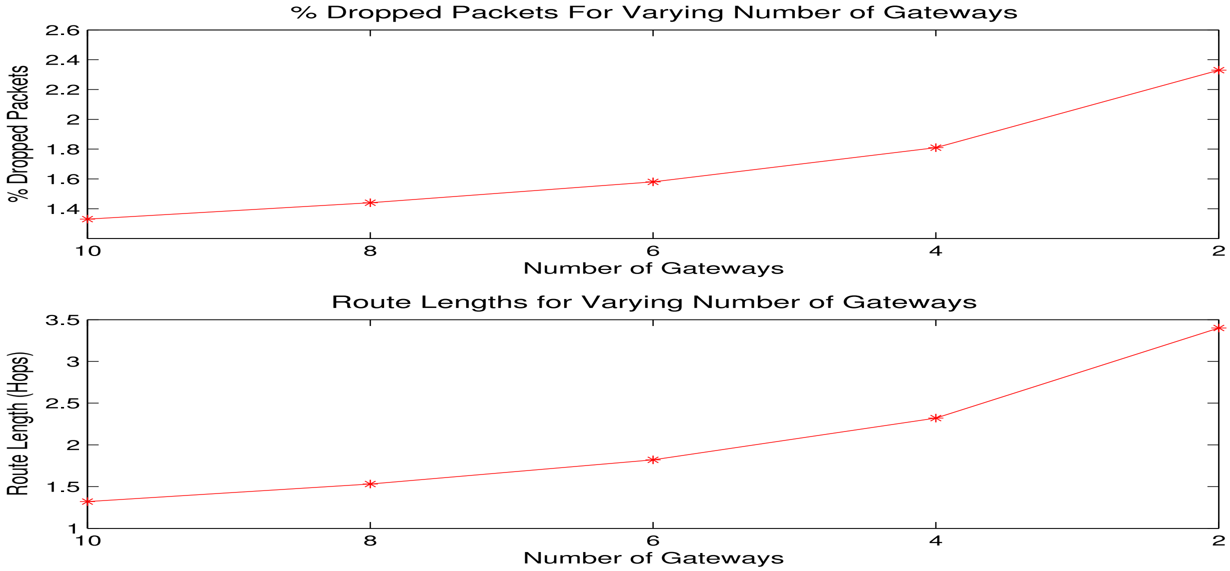

Fig. 7. Effect of varying number of gateways (keeping number of nodes fixed) on route lengths and, thus, dropped packets due to route failures

for this lifetime. The high frequency of route failures manifested as a large number of route requests sent to pre-emptively create new routes to avoid failure. Thus, the advantages of shorter route lengths by using routes through oncoming vehicles are lost due to excess overhead incurred in maintaining these routes.

C. *Effect of Mobility Pattern on Routing Performance*

Here we look at the packet delivery ratios achieved by the PBR protocol for varying randomness (controlled by p_r) in the mobility pattern. This will give an idea of the effect mobility patterns have on PBR routing performance. It can be expected that as mobility patterns become more random, the utility of predicted lifetimes will decrease. We are interested in quantifying the degree of randomness which affects performance due to route failures. Based on our study of vehicle and gateway density in the previous subsection, we choose three different combinations of node and gateway density which achieve a 99.9% connectivity. This allows us to focus on the effect of route failures on packet delivery ratio without worrying about lack of connectivity. The combinations we consider are {2,30}, {40,7} and {60 1} with the first entry being the number of nodes and the second being number of gateways. We look at varying packet delivery ratios and find the maximum value of p_r above which a certain ratio cannot be achieved. The greater p_r, the lower will be the packet delivery ratio. Thus, by varying p_r we can find

Fig. 8. Effect of mobility pattern randomness on packet delivery ratio for three different node and gateway density pairs

the maximum tolerable p_r to achieve a specific packet delivery ratio. With any value of p_r above this value, the specific packet delivery ratio may not be achievable.

As shown in Figure 8, for a large gateway density (about 1 per 66m in the case of $\{2{,}30\}$), the packet delivery ratio is usually very high for all possible values of p_r up to the maximum possible 0.5[5]. As the gateway density decreases, the maximum tolerable p_r starts dropping. For the $\{60{,}1\}$ case, the packet delivery ratio drops below 98% with a p_r of just 0.05. The decrease of maximum tolerable value of p_r with decreasing gateway density is because of increasing route lengths between nodes and their gateways. The average route lengths for the three sets of densities were found to be 1.001, 1.6 and 4.3 hops respectively.

Thus, we can conclude from this experiment that the effect of randomness is the smallest when the gateway density is very high; higher the density of gateways on the road, the smaller will be the ill-effects of randomness and better the performance of PBR.

[5] Above $p_r = 0.4$, the value of AGG is decreased from 0.2 to 0 as p_r increases to 0.5 to look at the complete range

D. Comparing PBR with Reactive and Proactive Protocols

The use of pre-emptive route creation by PBR protocol to offset route failure requires a larger number of routes to be created. To get a better perspective on the performance of PBR protocol, we compare it with two other protocols with a proactive and reactive flavor, with percentage of packets dropped due to route failures and number of route requests sent as metrics. The reactive protocol implemented in our simulator differs from the PBR protocol in that it creates routes on demand, only when a route does not exist or an existing route has failed. That is, the prediction algorithm is not employed. The proactive protocol on the other hand looks for new routes at specified intervals of duration `period` with the hope that new fresher routes can replace existing ones, reducing the likelihood of route failure. Thus, the proactive protocol is similar to the PBR protocol in that it also pre-emptively creates routes. We set `period` equal to 10 seconds for our experiment.

Figure 9 compares the three protocols for varying p_r with a number of nodes set to 40 and number of gateways set to 10. This node and gateway density allows us to again look at dropped packets due to route failures independent of connectivity. The effect of increasing p_r is mainly only on the PBR protocol. The smaller number of dropped packets by the other two protocols for a small p_r is because, with a low degree of randomness, nodes stay on at their same average speeds which results in much fewer route failures. The dropped packets ratio of PBR is much smaller than the reactive protocol, which is to be expected. Surprisingly, there is big difference between PBR and the proactive protocols as well. The reason for this is that even though both protocols employ pre-emptive route creation, PBR is able to predict when routes break and adjust its route creation interval accordingly. This decreases route failures and also reduces the number of routes created because if routes are more stable, pre-emptive route creation intervals will be larger. Thus, PBR is able to use prediction to adapt the pre-emptive route creation interval. With increasing p_r, the advantage of PBR reduces as studied in the previous experiment. As expected, the reactive protocol creates the least number of routes. Thus, even for mobility patterns with worst case values of p_r, it can be observed that there is big advantage in using PBR protocol due to much fewer dropped packets albeit at the unavoidable cost of higher overhead. The control overhead of PBR is much smaller than a proactive protocol because it is able to make educated guesses about when to create a new pre-emptive route.

We also compare the three protocols for a scenario where the gateway density varies during the simulation run. The gateway distribution is also non-random, with some areas of the road having most of the gateways. Out of 10 gateways used in above experiment, we remove 5 gateways three times during

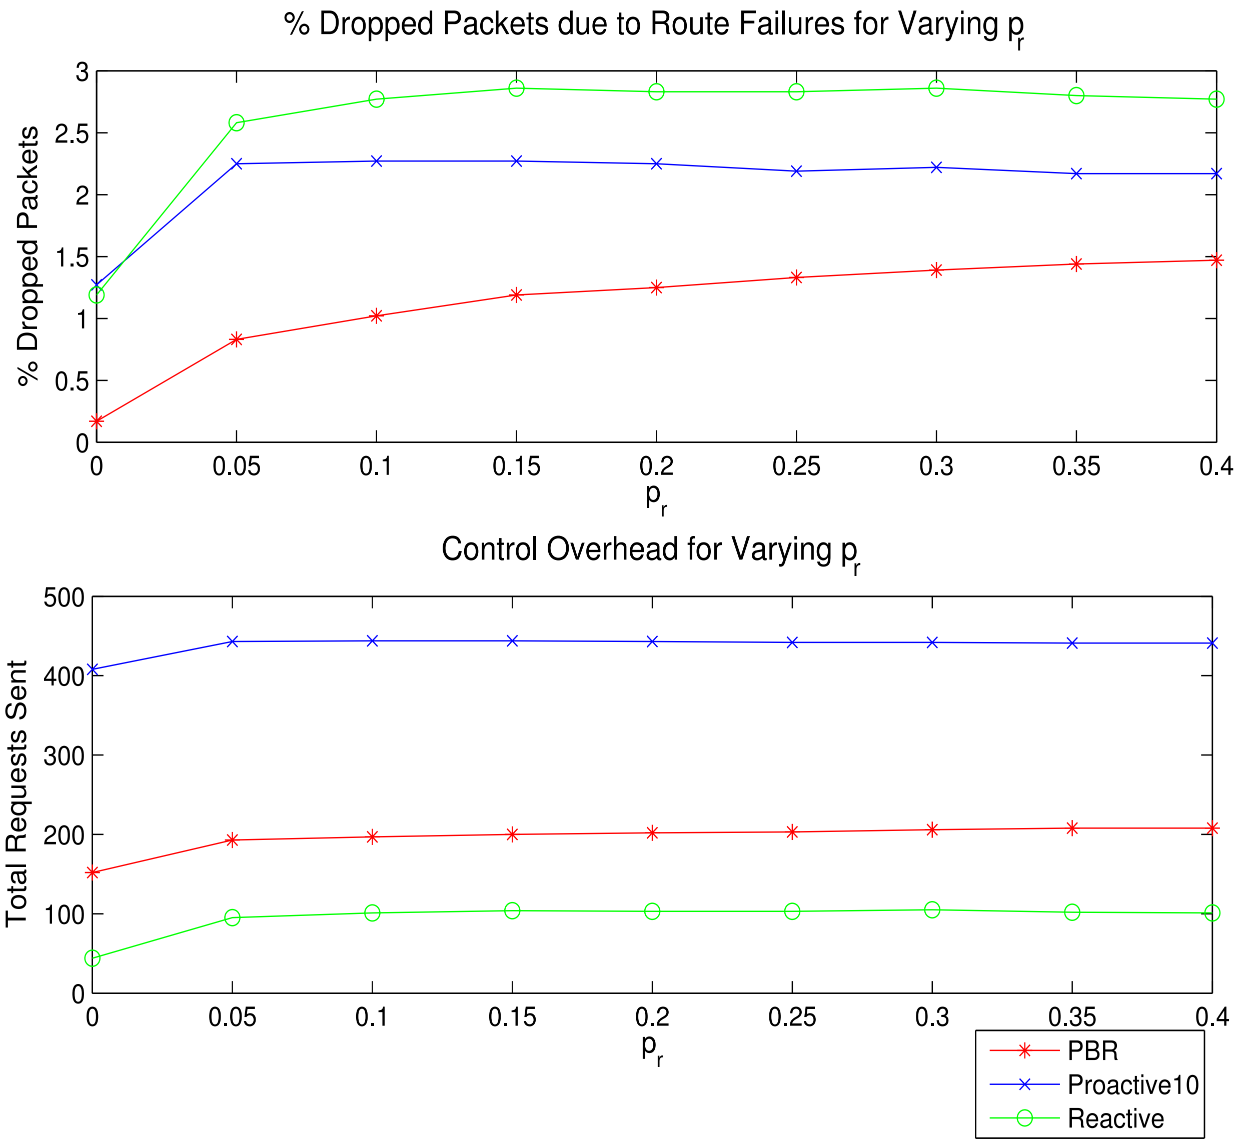
% Dropped Packets due to Route Failures for Varying p_r

Control Overhead for Varying p_r

Fig. 9. Comparing PBR with Reactive and Proactive protocols

the simulation run (at 500, 1500 and 2500 seconds) of 1 hour and put them all back in at the center of the highway stretch after intervals of 500 seconds. Figure 10 shows the results, which are similar to those of the previous plot comparing the three protocols. The only difference is that more packets are dropped and a higher overhead is incurred to re-construct routes for all protocols, due to the low density of gateways in some areas and at certain periods. In fact, it appears that the PBR protocol has widened the gap a bit in % dropped packets over the Reactive and Proactive protocols, with no trace of any disadvantage due to the different gateway density and distribution scenario.

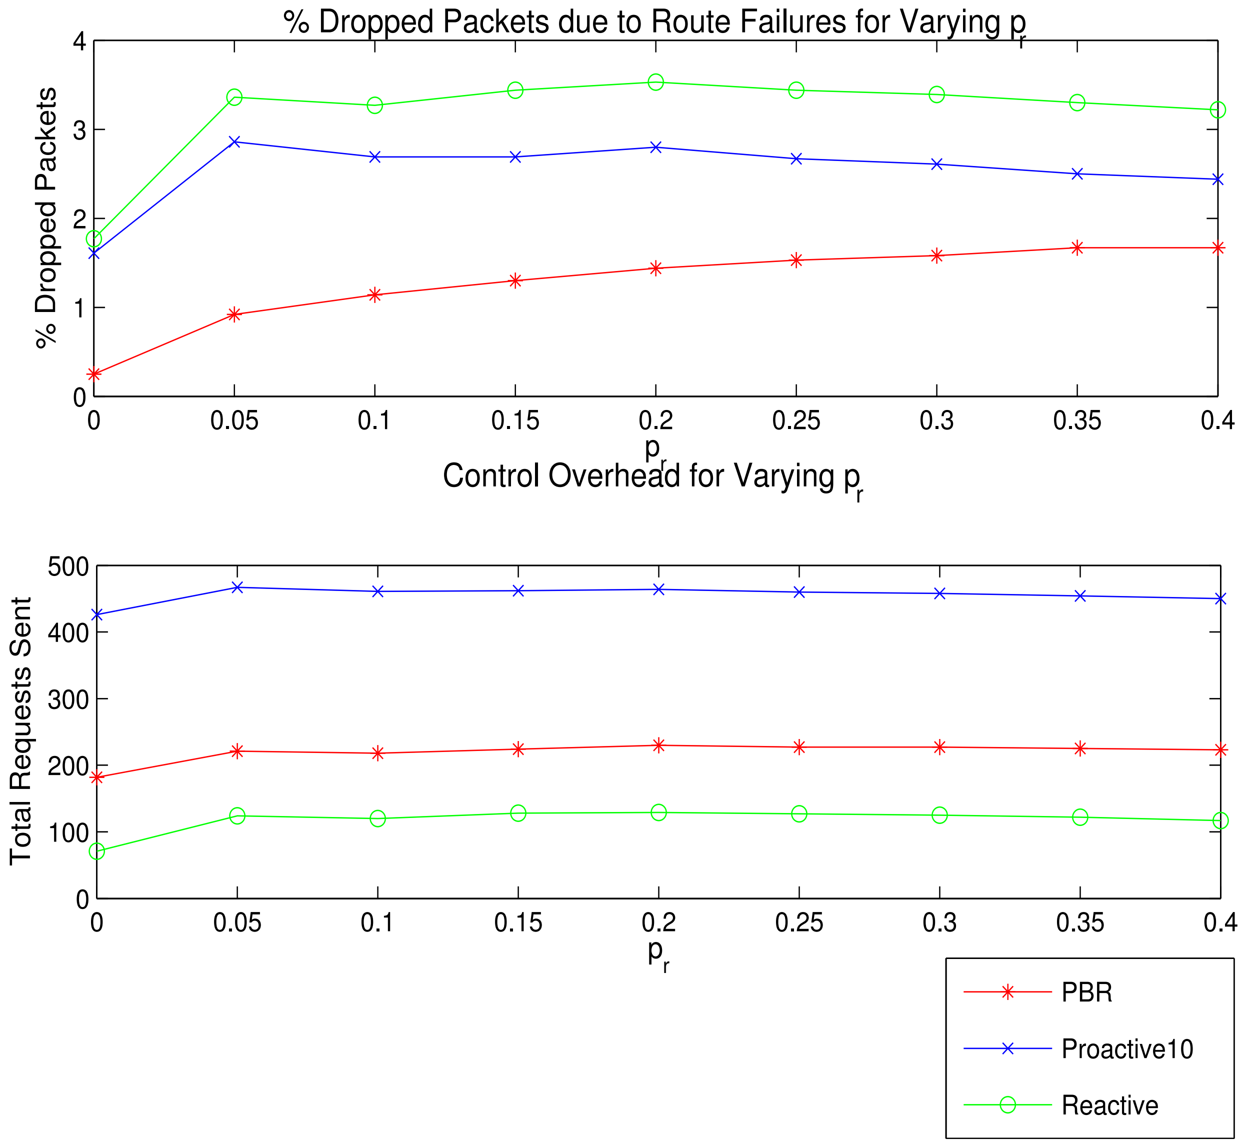

Fig. 10. Comparing PBR with Reactive and Proactive protocols with a variable gateway density scenario

The higher overhead of PBR is not a big penalty for the mobile gateway scenario due to the following two reasons. For general wireless ad-hoc network settings, control overhead can lead to smaller throughput and higher energy consumption. These networks are usually energy constrained due to devices with limited battery power. Also when used as access networks, they require high throughput. In such a setting, control overhead merits almost as much attention as the performance of the routing protocol. Vehicular networks are different. Energy supply is relatively abundant and the amount of power required for wireless communication can be safely ignored. On the throughput side, when nodes connect to the

Internet, the bottleneck link in terms of bandwidth is the WAN link (Figure 1) from the gateways to the base station. Comparatively, the WLAN network has much higher bandwidth and should be able to absorb the extra control overhead which is only required among vehicles communicating using their WLAN interface.

E. Minimizing Gateway Switching

Here we explore the performance of the two PBR variants, PBR-S and PBR-M, presented in Section IV. The mobility randomness parameter p_r was set to 0.25. We vary the number of nodes from 35 to 55 keeping the number of gateways equal to 10. These node densities result in very high connectivity (99.9%). We felt it was more interesting to vary the number of nodes which plays an important role in retaining paths to gateways for this experiment to study the different PBR variants. As node density increases, it can be expected that a node's gateway will be retained for longer periods of time and also present more diversity in selecting paths to gateways.

Figure 11 shows the plot for this experiment. The unmodified version of PBR as earlier is also shown for reference. There are only small differences between PBR and PBR-M. Being able to select longer lasting gateways has increased the time spent with a gateway and decreased the number of routes created. But, the increased predicted route lifetimes results in more route failures. The prediction algorithm uses only information gathered when a route is created, with no feedback until a route fails. Thus, as predicted lifetimes increase, the chances of the routes failing before that time increase. Thus, the strategy of selecting gateways based on predicted lifetime does not seem to be too helpful due to the limitations of our prediction algorithm. There is, however, a noticeable reduction in control overhead using PBR-M making it useful for situations with higher emphasis on minimizing control overhead rather than route failures.

The first thing of note when comparing PBR and PBR-S is that the percentage of dropped packets due to route failures of PBR-S is much higher than PBR. This is because, the route lengths have increased significantly (since these protocols try to stick to gateways regardless of how far they are), which increases the probability of route failure. The control overhead is also much more when nodes try to retain their gateways because the increased route failures require more frequent route creation. Finally, as expected, the average time spent with a gateway increases for PBR-S protocol due to smaller number of gateway switches. It is interesting to note that the PBR-S protocol only suffers only an additional 2% packet drops than the PBR protocol, while staying with a gateway for a much longer period. For smaller gateway densities than the one considered for this experiment, however, the effect of randomness on

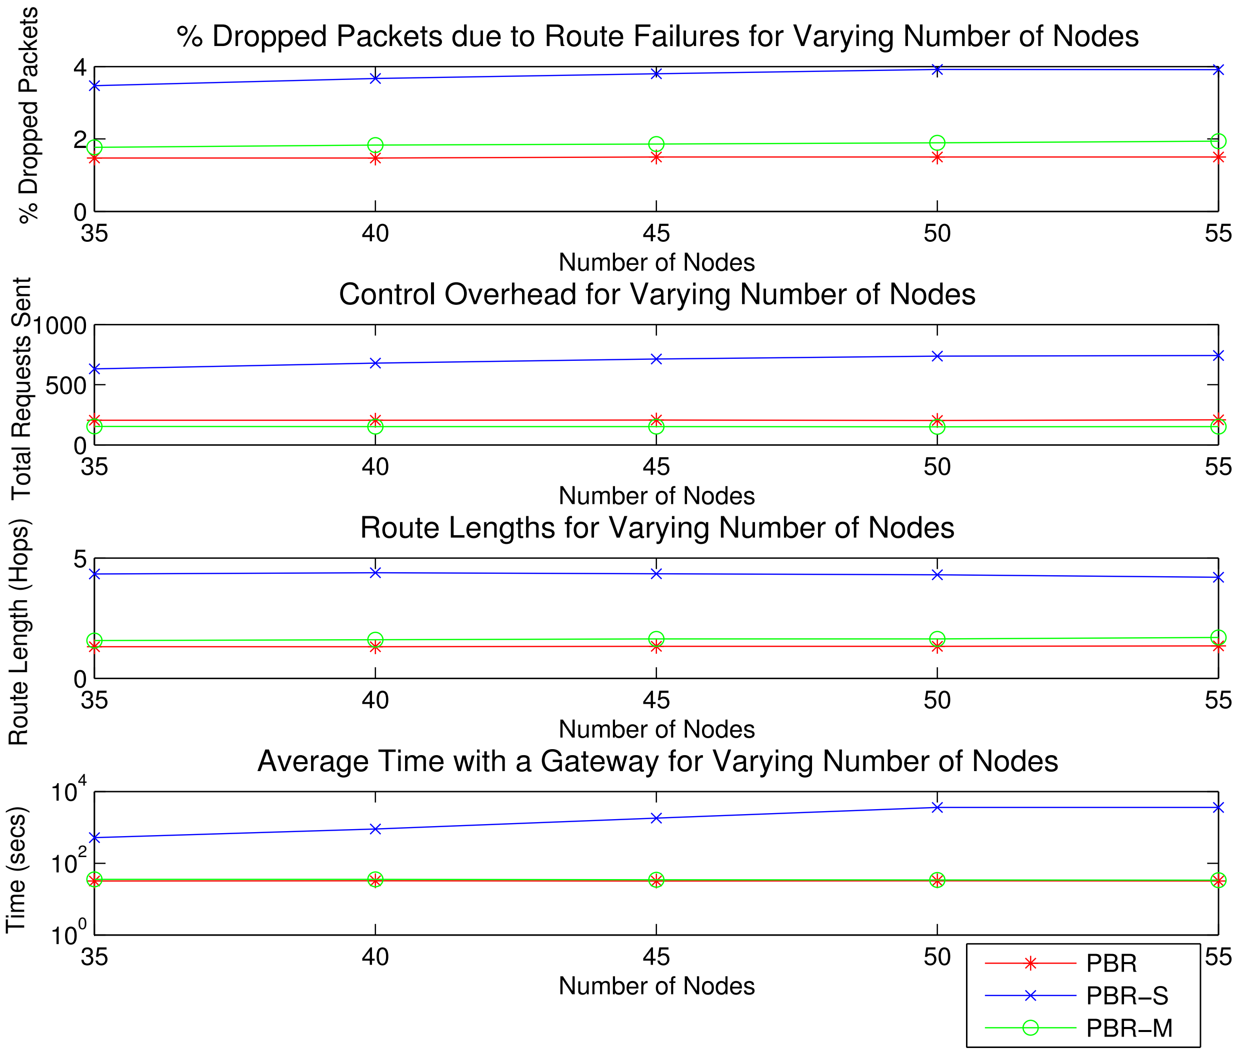

Fig. 11. Studying the performance of PBR when minimal gateway switching is desired by comparing with PBR-S and PBR-M

route failures will be far greater due to larger route lengths for PBR-S, making PBR more attractive with its shorter route lengths. In the case of PBR-S, a higher node density can often result in larger route lengths and increased route failures by allowing a node to retain its gateway for long periods of time regardless of how far it is. That effect could not be shown here unfortunately due to the limited highway stretch simulated as a torus which limits the maximum route length around 5 hops as shown for a length of 2000m and communication range of 200m.

Thus, applications which require minimal gateway switching will perform well by using PBR-S subject to two conditions (i) the gateway density is high enough to offset the increased path lengths that result from retaining gateways (ii) A limit be kept on how far a gateway can be for it to be retained. If gateway switching can be supported, the policy of finding the nearest gateway seems the best choice for the mobile gateway scenario. Other possible considerations like load on gateways and QoS merit further study on selecting and retaining gateways, but is outside the scope of this work.

VII. CONCLUSIONS

The routing aspects of the scenario of vehicles on the road connecting to the Internet through mobile gateways was introduced. It was found that the predictable motion of vehicles on the road along with readily available location and velocity information could be exploited to predict route failures. Such a prediction based routing (PBR) protocol was presented, which used predicted route lifetimes to take preemptive action, minimizing route failure. This protocol was extensively evaluated for varying vehicle densities and mobility patterns and compared against proactive and reactive protocols. Ideal vehicle and gateway densities for the mobile gateway scenario were identified which could guarantee application performance. The effect of route lengths and mobility patterns on routing performance were analyzed and a high gateway density recommended to minimize the ill effect of both. It was found that PBR offers significant reductions in route failures leading to higher packet delivery ratios with the control overhead kept in check due to its ability to predict route lifetimes. Moreover, the use of oncoming vehicles on routes was discouraged except when forward moving vehicle densities are low enough to affect connectivity.

In the future, as connectivity on the road is used for many applications, static gateways can be used to supplement mobile gateways making up for the lack of density in certain areas and at certain times. Wireless LAN technologies other than IEEE 802.11b with a higher communication range would also prove useful to offset reduced density on state highways and rural roads which generally do not have enough vehicle density. A higher range would also reduce the route lengths which reduces route failures. The trade off of increasing range is decreasing capacity due to radio interference among nodes. This issue needs to be researched further, especially looking at power control to vary the range depending on vehicle density. Also, pricing models for the scenario can be used to ensure the viability of the scenario for all concerned.

REFERENCES

[1] V. Namboodiri and L. Gao, “A Study on the Feasibility of Mobile Gateways for Vehicular Ad-Hoc Networks,” in *Proc. of ACM VANET*, 2004, pp 66-75.

[2] H. Füßler, M. Mauve, H. Hartenstein, M. Käsemann, and D. Vollmer, “A Comparison of Routing Strategies for Vehicular Ad-Hoc Networks,”, Dept. of Math. and C.S., Univ. of Mannheim, Tech. Rep. TR-02-003, July 2002.

[3] M. Bai, L. Wolf, and W. Franz, "Mobile Internet Access in FleetNet," *13. Fachtagung Kommunikation in Verteilten Systemen KiVS* , 2003.

[4] M. Bechler, O. Storz, W. Franz, and L. Wolf, "Efficient Discovery of Internet Gateways in Vehicular Communication Systems," in *Proc. of IEEE VTC*, Spring 2003, vol. 2, pp. 965-969.

[5] J. Ott and T. Kutscher , "Drive-thru Internet: IEEE 802.11b for "Automobile" Users," in *Proc. of IEEE INFOCOM* , 2004, vol. 1, page 373.

[6] V. Bychkovsky, B. Hull, A. K. Miu, H. Balakrishnan and S. Madden, "A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks", in *Proc. of ACM MOBICOM* , 2006, pp. 50-61.

[7] 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 *Proc. of IEEE INFOCOM*, 2003, vol. 1, pp. 503-512.

[8] H.-Y. Hsieh and R. Sivakumar, "Internetworking WWANs and WLANs in Next Generation Wireless Data Networks," in *Proc. of International Conference on 3G Wireless and Beyond*, 2002.

[9] 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 *Proc. of ACM/IEEE MOBICOM*, 1998, pp. 85-97.

[10] M. Mauve, J. Widmer, and H. Hartenstein, "A Survey on Position-Based Routing in Mobile Ad-Hoc Networks," In *IEEE Network*, 15(6) pp. 30-39, Nov./Dec. 2001.

[11] E. M. Royer and C.-K. Toh, "A Review of Current Routing Protocols for Ad-hoc Mobile Wireless Networks," *IEEE Personal Communications*, 6(2) pp. 46-55, April 1999.

[12] S. Murthy and J. Garcia-Luna-Aceves, "An Efficient Routing Protocol for Wireless Networks," *Mobile Networks and Applications* 1(2) pp. 183-197, Oct. 1996.

[13] C.E. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector (DSDV) Routing for Mobile Computers," *Computer Communications Review* , pp. 234-244, Oct. 1994.

[14] D.B. Johnson and D.A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," *Mobile Computing* , edited by T. Imielinski and H. Korth, Kluwer Academic Publishers, pp. 153-181, 1996.

[15] V. D. Park and M. S. Corson, "A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks," in *Proc. of IEEE INFOCOM*, 1997, pp. 1405-1413.

[16] C.E. Perkins and E. M. Royer, "Ad hoc On-Demand Distance Vector Routing." in *Proc. of IEEE WMCSA*, 1999, pp. 90-100.

[17] Y. Ko and N. H. Vaidya, "Location-Aided Routing (LAR) Mobile Ad Hoc Networks," in *Proc. of ACM/IEEE MOBICOM*, 1998, pp. 66-75.

[18] T. Goff, N.B. Abu-Ghazaleh, D.S. Pathak, and R. Kahvecioglu, "Preemptive Routing in Ad Hoc Networks," *Journal of Parallel and Distributed Computing*, 63(2) pp. 123-140, Feb. 2003.

[19] W. Su, S. Lee, and M. Gerla, "Mobility Prediction and Routing in Ad Hoc Wireless Networks," *International Journal of Network Management*11(1) pp. 3-30, Jan./Feb. 2001.

[20] 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) pp. 75-88, April 2002.

[21] 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*, July 2002.

[22] V. Muchuruza and R. Mussa, “Speed on Rural Interstate Highways Relative to Posting the 40 mph Minimum Speed Limit,” *Journal of Transportation and Statistics* 7(2/3), 2004.

[23] N. Bellomo, V. Coscia, and M. Delitala, “On the Mathematical Theory of Vehicular Traffic Flow I: Fluid Dynamic and Kinetic Modelling,” *Mathematical Models and Methods in the Applied Sciences*, 12(12) pp. 1801-1843, 2002.

[24] B. S. Kerner, “The Physics of Traffic,” *Physics World* , vol. 12, pp. 25-30, Aug. 1998.

[25] 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*, 1(4) pp. 383-403, Nov. 2003.

[26] C. Bettstetter, “Smooth is Better than Sharp: A Random Mobility Model for Simulation of Wireless Networks,” in *Proc. ACM MSWiM*, 2001, pp. 19-27.

[27] K. Nagel, P. Stretz, M. Pieck, S. Leckey, R. Donnelly, and C. L. Barrett, “TRANSIMS Traffic Flow Characteristics,” Los Alamos National Laboratory Report LA-UR-97-3530, July 1997.

[28] L. E. Owen, Y. Zhang, and L. Rao, “Traffic Flow Simulation using CORSIM,” in *Proc. of Winter Simulation Conference* , 2000, pp. 1143-1147.

[29] Y. C. Tseng , S. Y. Ni , Y. S. Chen, and J. P. Sheu, “The Broadcast Storm Problem in a Mobile Ad Hoc Network”, *Wireless Networks* , 8(2/3) pp. 153-167, Mar. 2002.

[30] “Traffic Speeds on English Trunk Roads:2003,” Transport Statistics, Department for Transport, UK, 2003.

[31] “Highway Statistics 2004,” U.S Department of Transportation Federal Highway Administration (FHWA), hm72, 2004.

[32] S. N. Humbad, “Suitability of Distributed Mobile Wireless Networking for Urban Traffic Congestion Mitigation,” Master in City Planning, Department of Urban Studies and Planning, MIT, June 2001.