Ad Hoc Relay Wireless Networks over Moving Vehicles on Highways *

Zong Da Chen
Harvard University
33 Oxford St
Cambridge, MA 02138
zdchen@fas.harvard.edu

HT Kung
Harvard University
33 Oxford St
Cambridge, MA 02138
htk@eecs.harvard.edu

Dario Vlah
Harvard University
33 Oxford St
Cambridge, MA 02138
dario@eecs.harvard.edu

ABSTRACT

Ad hoc networks can be formed on highways among moving vehicles, each equipped with a wireless LAN device. However, during times of low traffic density, it is likely that such networks are disconnected. This paper tests the hypothesis that the *motion* of vehicles on a highway can contribute to successful message delivery, provided that messages can be *relayed*—stored temporarily at moving nodes while waiting for opportunities to be forwarded further. Using vehicle movement traces from a traffic microsimulator, we measure average message delivery time and find that it is shorter than when the messages are not relayed. We conclude that ad hoc relay wireless networks, based on wireless LAN technologies, have potential for many emerging applications of this kind.

1. INTRODUCTION

Wireless connectivity on vehicles is emerging as an important mode of communication; already, large car makers are rolling out infrastructure-based low bandwidth wireless services in support of applications termed *automotive telematics*. Such applications are a combination of telecommunications and computation, such as route planning using GPS signals or remote diagnostics using data from sensors built into vehicles.

It is more challenging to provide high-bandwidth networking to fast moving vehicles. Cell sizes for high-bandwidth communication are usually small, resulting in high handoff rates as vehicles cross the cell boundaries. A recently proposed solution even suggests to let wireless base stations follow the highway traffic on a rapidly moving conveyor rail installed along the highway divider [5].

*This research was partially supported by the Air Force Office of Scientific Research Multidisciplinary University Research Initiative Grant F49620-97-1-0382 and as well as Microsoft Research, Nortel Networks, and Sun Microsystems.

Ad hoc networks formed by high-bandwidth, short range wireless devices, such as based on the 802.11 wireless LAN standard [7], are well suited for moving vehicles. Their deployment on individual vehicles doesn't require any infrastructure; also, ad hoc routing adapts to node mobility. A number of interesting applications are possible in such networks; we categorize them into three groups.

Traditional applications include all of the popular Internet services such as Web browsing, E-mail, video streaming, etc. These can be enabled in a general way by equipping cars with access points for existing portable devices like notebooks or PDA's. Connectivity to the Internet at large can be achieved through gateway nodes placed along the highway.

Location-aware applications provide geographic location information for points of interest to users. For example, a directory may provide real-time driving directions, or listings of local services such as gas stations or businesses.

Localized applications are run collaboratively by groups of nearby nodes, and thus match the ad hoc routing model best. For example, vehicles may engage in chat, exchanging short messages, sharing multimedia files, or engage in video-conferencing. There is a possibility for offering services; we can imagine an "information vendor" driving along the highway, carrying a library of music files for download. Vehicles belonging to highway service crews can announce themselves, automatically informing drivers about oncoming road work areas. Disabled or crashed vehicles may use acceleration sensors to trigger automatic advertisements of an accident, as pointed out in [1]. Morris et al. [13] point out additional localized applications like collaborative road congestion monitoring, route planning, fleet tracking, etc.

With ad hoc networks deployed on moving vehicles, network partitions due to limited radio range become inevitable when traffic density is low—such as at night, or when few vehicles carry a wireless device. A key question to ask is whether it is possible to deliver messages in spite of partitions, by taking advantage of predictable node movement on the highway creates opportunities to relay messages in a store-and-forward fashion.

This paper tests the hypothesis that the *motion* of vehicles on a highway can significantly contribute to successful message delivery, provided that messages can be *relayed*—

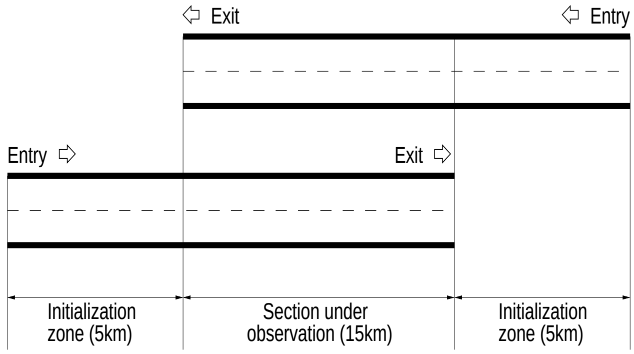

Figure 1: Highway geometry used by the traffic simulator of this paper

temporarily stored at moving nodes while waiting for opportunities to be sent further. We go about testing this hypothesis by simulating movement of vehicles on a highway, simulating an ad hoc network over the vehicles, and measuring the performance of the network as traffic density decreases.

In Section 2, we describe our simulation environment. In Section 3, we present and discuss the experimental results. We present some related work in Section 4, and finish with conclusions in Section 5.

2. SIMULATION ENVIRONMENT

Our simulation environment consists of two components. The first is a traffic microsimulator which produces accurate movement traces of vehicles traveling on a highway. The second is a network simulator which models transport of messages among the vehicles.

2.1 Traffic Microsimulator

CORSIM (CORridor SIMulator) which we use in this paper is a microscopic traffic simulator developed by the Federal Highway Administration [3]. CORSIM models behavior of human drivers by approximating a set of common driver decisions, such as slowing down or changing the lane in the vicinity of slower vehicles ahead. These decisions cause the acceleration and orientation of vehicles to change, resulting in realistic movement patterns.

There are three inputs to the simulator: a highway geometry, a free-flow speed, and an input rate. The geometry in our simulations is very basic—as depicted in Figure 1, it is a straight highway segment with two directions, each composed of one or more lanes.

The free-flow speed parameter is the mean speed that travelers achieve when they move unconstrained by other vehicles or obstacles. The actual free-flow speeds assigned to vehicles are distributed normally around that mean, with a default standard deviation of 8mph. The free-flow speed parameter we used in our simulations was 50mph.

The input rate parameter controls the period at which the simulator generates new vehicles. At the end of each period, a new vehicle is placed at one of the entry points indicated in Figure 1, but only if enough room is available—in case of congestion, no new vehicles are generated.

# Lanes	Lowest density (cars/km)	Highest density (cars/km)
1	3.1	79
2	3.1	172
3	3.2	273
4	3.2	336
5	3.5	325

Table 1: Summary of the highway traces obtained from CORSIM and used to drive network simulations. Note the lower maximum densities observed on one- to three-lane highways, due to the highways' lower total capacity

Vehicles near the entries move in very regular formations, due to the uniform period at which they are generated. In order for the vehicles to assume more random positions, we first let them cross 5km *initialization zones*. Only when they reach sections under observation are their positions recorded. Once they leave these sections, vehicles are removed from the simulation.

The output of each simulation is a trace of vehicle positions taken at one second intervals. All our simulations lasted 300 simulated seconds. We ran a total of 170 simulations, for highways with one to five lanes on each side, and total input rates varying between 5 and 800 cars per hour. The average vehicle densities observed in the 170 simulation runs are summarized in Table 1. For small numbers of lanes, maximum traffic densities are constrained by the capacity of the highway.

2.2 Network Simulator

We simulate a wireless network over vehicles driven by the traces summarized in Table 1. We assume the wireless devices installed on vehicles all have the same fixed radio range r. Our simulator omits detailed representation of protocol layers and radio propagation; instead, it maintains a network connectivity graph, and positions and age of the messages in transit. Messages are propagated greedily each timestep, by hopping to the neighbor closest to the destination. This amount of state is sufficient for the purpose of finding the delay due to mobility.

Before moving to the measurements obtained from the simulator, let us describe two kinds of transmission patterns in our wireless network that we call *pessimistic* and *optimistic* forwarding, which are distinguished by how long the messages are permitted to stay in intermediate nodes. In pessimistic forwarding, a message is dropped whenever no next hop exists for its destination. This is how forwarding works in most ad hoc network implementations. [1] In optimistic forwarding, messages without next hops may remain on intermediate nodes for some time, hoping that physical movement of network nodes eventually creates a forwarding opportunity. For example, in a geographic routing protocol such as [9] [10] coupled with optimistic forwarding, messages would not be dropped, but instead held and then forwarded greedily as soon as a new neighbor node closer to the des-

[1] Some on-demand protocols, like DSR or AODV, may buffer packets during a route-request or route-repair period.

tination was detected. We now describe how we measured performance of both of these kinds of forwarding.

Under pessimistic forwarding, we measure the average amount of time a sender S waits before a direct route to a destination D becomes available. We obtain this time directly from the movement traces as half the average duration of network partitions between S and D, and call it the *pessimistic delay*. We note that such delay is a lower bound only observable in an ideal network with perfect routing information (either local or global)—in practice, route computation may be affected by additional factors, such as timing information of a particular routing protocol. For example, suppose the network uses a geographic routing algorithm. Then, even though a destination D becomes reachable from source S due to movement of some nodes, the messages from S will still not reach D until the new topology information is learned using beaconing between neighbors.

With optimistic forwarding, messages are not dropped, so we measure the average time messages spend in the network. We call this quantity the *optimistic delay*.

3. RESULTS

We now present the experimental results. We begin by describing the experimental setup, and follow with a set of three measurements.

3.1 Experimental Setup

For the experiment of this section, we use the following traffic scenario. A single packet is sent by each car entering the highway. The destination of the packet is chosen to be 10 km away, in the entering car's direction; however, if no such car happens to be within r of 10km away, then no packet is sent.

We perform the delay measurements on the highway traces we have summarized in Table 1, for radio range $r = 200m$.

3.2 Comparing the Optimistic and Pessimistic Delays

Figure 2 shows the behavior of measured optimistic and pessimistic delay as density of vehicles on the highway increases. At high densities, both delays are close to zero since the likelihood that a destination is unreachable is very low. However, pessimistic delay rises more sharply than optimistic delay as density falls, indicating that mobility succeeds helping the delivery of optimistic messages. The significance of this result is in showing that a network taking advantage of node mobility can operate in lower densities while maintaining same average delay. For example, according to the figure, applications which tolerate 200 seconds of delay can operate in networks that are half as dense, if messages are forwarded optimistically.

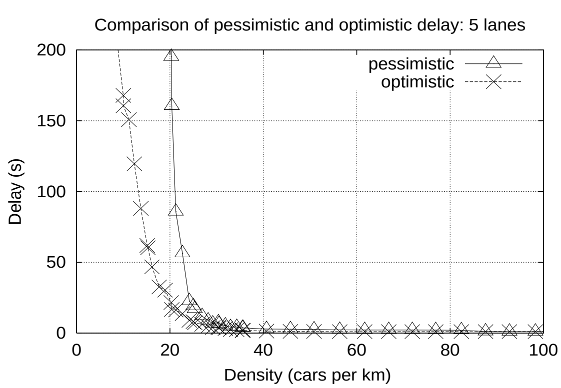

Figure 2: Optimistic vs. pessimistic delay, 5 lanes, bidirectional

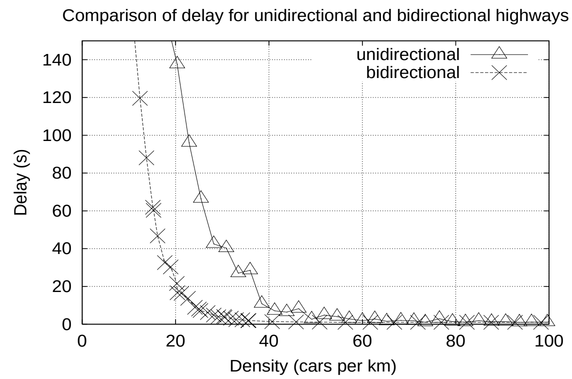

Figure 3: Optimistic delay on unidirectional and bidirectional highways with 5 lanes in each direction

3.3 Unidirectional and Bidirectional Highways

Figure 3 shows two sets of optimistic delays, measured on unidirectional and bidirectional highways. The delay at any given density is lower for bidirectional traffic, indicating that the rapid relative motion of vehicles in two directions makes a significant contribution to reducing message delivery delay.

3.4 Effect of the Number of Lanes

Figure 4 shows the optimistic delays measured on bidirectional highways with 1 to 5 lanes. It is evident that small numbers of lanes have detrimental effect on network performance. With only one lane, inevitably there are slow vehicles which accumulate tails of followers, and cause the

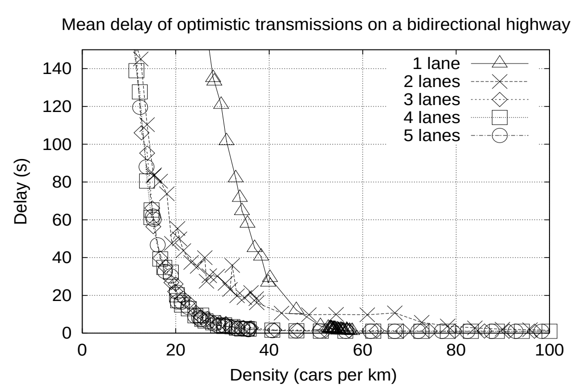

Figure 4: Optimistic delay on bidirectional highways with 1 to 5 lanes in each direction

#lanes	Density Low (50 cars/km)	Medium (100 cars/km)	High (150 cars/km)
1	Fail	No data	No data
2	Fail	Fail	No data
3	Pass	Fail	Fail
4	Pass	Pass	Fail
5	Pass	Pass	Pass

Table 2: Outcomes of the χ^2 test for uniform distribution of cars, for a range of densities on highways with 1 to 5 lanes

distribution of vehicles to become clustered. The likelihood of network partitions then increases, leading to increased delay. An additional lane creates the possibility for the queued followers to switch lanes and pass slow cars. However, at medium densities this may not be possible because both lanes are occupied—indeed, the plot for a two-lane highway exhibits an increase in delay for high density values. With three or more lanes, there is enough freedom of movement to prevent clustering, leading to lower observed delay.

Intuition tells us that the distribution of cars, when their movement is uninhibited, should be uniform. In order to verify this, we tested whether the distribution of intervals occuring in the traces matches that of intervals in a uniformly distributed set of points. Table 2 shows the outcomes of χ^2 tests [8] at a 80% level of confidence. The distributions become less uniform in two cases: when density of traffic grows, causing the vehicles to become more regularly spaced, or when the number of lanes decreases, and slow vehicles cause clustering.

Under the assumption that the vehicle distribution is uniform, it is possible to develop a simple mathematical analysis of optimistic and pessimistic delays [2]. According to the χ^2 test results, such analysis predicts delay more accurately for highways with more lanes and sparser traffic.

4. RELATED WORK

The idea of using node mobility to help network performance has been investigated by some recent work. Briesemeister and Hommel [1] proposed a multicast protocol for the highway environment, and measured the fraction of nodes to which optimistically forwarded messages were successfully delivered in a fixed amount of time. In contrast, our work measures performance as average delay required to traverse a certain fixed distance.

Further, Li and Rus [11] looked at a scenario in which the mobile nodes were under their control, so that the trajectories could be modified to reduce delay. Finally, Grossglauser and Tse [6] started from a model with random node movements, and showed that the motion leads to constant throughput as the network size increases.

5. CONCLUSION

Let us finish with the following three points. First, we found that mobility of nodes on the highway improved end-to-end transmission delay if messages were relayed—that is, if they were held at intermediate nodes until favorable forwarding paths appeared. The improvement leads to practical delay values at smaller traffic densities than was possible even for pessimistically forwarded messages, with ideal routing information. Thus, we conclude that our initial hypothesis is valid.

Next, the improvement was higher for traffic scenarios with more relative movement. There were two sources of relative movement: traffic in oposing directions on bidirectional highways, and multi-lane traffic within the same direction.

Finally, since the delays measured at low vehicle densities increase to the order of seconds, such conditions cannot support delay sensitive applications like interactive multimedia which require low delay guarantees. However, we believe there are applications, such as some of the localized applications which are delay tolerant and are therefore suitable for this environment.

6. REFERENCES

[1] Briesemeister, L. and Hommel, G., Role-Based Multicast in Highly Mobile but Sparsely Connected Ad Hoc Networks, in *Proceedings of the First Annual Workshop on Mobile Ad Hoc Networking & Computing (MobiHOC)*, Boston, MA, USA, August 2000

[2] Chen, Z. D., Wireless Networking on Highways, Senior Thesis, Harvard University, Cambridge, MA, USA, April 2001

[3] CORSIM User Manual, Version 1.01, The Federal Highway Administration, US Department of Transportation, 1996

[4] Feller, W., An Introduction to Probability Theory and Its Applications, Volume II, John Wiley & Sons, Inc., New York, 1971

[5] Gavrilovich, C. D., Broadband Communication on the Highways of Tomorrow, in *IEEE Communications Magazine*, April 2001

[6] Grossglauser, M., and Tse, D., Mobility Increases the Capacity of Ad Hoc Wireless Networks, in *Proceedings of Infocom 2001*, Anchorage, AK, USA, April 2001

[7] IEEE Computer Society LAN MAN Standards Committee, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Std 802.11-1997, The Institute of Electrical and Electronics Engineeers, New York, New York, 1997

[8] Jain, R., The Art of Computer Systems Performance Analysis, John Wiley & Sons, Inc., New York, 1991

[9] Karp, B. and Kung, H.T., Greedy Perimeter Stateless Routing for Wireless Networks, in *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*, Boston, MA, USA, August 2000

[10] Li, J., Jannotti, J., De Couto, D. S. J., Karger, D. R. and Morris, R., A Scalable Location Service for Geographic Ad Hoc Routing, in *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*, Boston, MA, USA, August 2000

[11] Li, Q. and Rus, D., Sending Messages to Mobile Users in Disconnected Ad hoc Wireless Networks, In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)* Boston, MA, USA, August 2000

[12] M25 Real Time Trial, The UK Highways Agency, May 2001, http://www.highways.gov.uk/m25rt/index.htm

[13] Morris, R., Jannotti, J., Kaashoek, F., Li, J., and Decouto, D., CarNet: A Scalable Ad Hoc Wireless Network System, in the *Proceedings of the 9th ACM SIGOPS European Workshop*, September 2000