

Use of Mobile Mesh Networks for Inter-Vehicular Communication

Dr. Paul Beckman

Dr. Sameer Verma

Dr. Ramesh Rao

Department of Information Systems
San Francisco State University
San Francisco, USA

Abstract—High-speed wireless computing networks are now economically feasible for home users, via 802.11b wireless protocols and their associated hardware. Such centralized “point-to-multipoint” installations typically allow a range of about 100 meters from the central access point. It is possible to use the same network protocols and hardware to construct a “wireless mesh” or “multipoint-to-multipoint” network in which any connected node can communicate directly or indirectly with any other connected node.

The research project described in this article applies wireless mesh networking to inter-vehicular communication. Three vehicles were connected as a wireless mesh network using laptop computers with 802.11b radio cards and mesh networking software. The vehicles were then driven on a highway in Northern California to collect data about network connectivity. The goal of the experiment was to prove that such a network could be quickly and easily constructed and that network connectivity could be maintained under normal driving conditions. Data was collected on network connectivity and time delay of network packet transmission.

I. INTRODUCTION

The experiment documented in this paper extended a current method of wireless networked computing, called Wi-Fi, to a new realm of use: inter-vehicular communication. Wi-Fi wireless computer networks that use the 802.11b data transmission protocols have recently become quite popular for home computer users. Consumer Wi-Fi products typically include an access point (AP) that connects by wire to a home-based high-speed Internet connection such

as a DSL line or cable modem and one or more mobile nodes that install as PCMCIA cards in mobile computer devices such as laptops. The remote devices are also available in several different form factors besides PCMCIA cards, such as add-on cards for PDAs. These networks are fairly low in power and therefore have a geographic range of about 100 meters. Recently, the 802.11b protocol standards have been expanded to allow higher maximum bandwidth delivery, from the existing 11 mbps maximum (for 802.11b) to 54 mbps (for 802.11g).

A standard application for such networks is to wirelessly connect many laptop computers to the Internet. These applications are sometimes used internally within a campus of one organization. They are also sometimes set up for commercial purposes, by WISPs (wireless Internet service providers), to provide paying customers access the Internet at local, regional, or national “hotspots” where branded Wi-Fi network access is available. These applications operate on financial payment systems where the customer either pays by the minute of use or for unlimited monthly access. Recently, non-profit organizations have taken the WISP concept to local neighborhoods and set up CWNs (community wireless networks) that provide essentially free wireless Internet access to local community citizens. It is unclear which of these service models will ultimately prevail [4], and in some cases, WISPs and CWNs compete directly for the same customers.

A. The Wi-Fi protocol

The Wi-Fi protocol was designed as a point-to-multipoint solution wherein the service

provider would set up one AP to which many remote devices could simultaneously connect. However, it is possible to use the 802.11b protocols and hardware but different network routing algorithms to allow for implementation of a multipoint-to-multipoint network. Such a network is called a “mesh” because any node in the network can connect to any other node that is within physical radio range of some connected node [2], [3], as opposed to a Wi-Fi network in which a node can only connect to the central AP and then only if that node is within physical radio range of the AP. Mesh networks can therefore extend the geographic range of the wireless network to the radio range of the furthest connected remote node.

The experiment described in this paper used the concepts, hardware, and software of mesh networking in an application significantly different than allowing a set of physically stationary users to connect to the Internet via an AP. The impetus for performing the experiment came from the realization that there are several similarities between the characteristics of the nodes in a standard mesh network using the 802.11b protocols and the characteristics of vehicles in normal highway traffic flow.

B. Similarities between mesh networks and traffic flow

Relevant characteristics of a wireless mesh network are:

1. nodes can randomly enter and leave the network
2. nodes must be within about 100 meters of each other
3. higher node density impacts data throughput

Relevant characteristics of highway traffic flow are:

1. vehicles can randomly enter and leave the traffic flow
2. vehicles often travel within about 100 meters of each other
3. higher traffic density impacts vehicle throughput

II. METHODOLOGY

The experiment was performed by running mesh networking software on three laptops each

in separate vehicles, each of which also had a remote wireless networking card. The vehicles were driven approximately 30 kilometers on a United States interstate highway in Northern California with an attempt to maintain normal highway driving speeds (approximately 100 to 110 kilometers/hour throughout) and distances (from about 30 meters to about 300 meters) between the vehicles.

Two of the laptops were Compaq Presario 1625s; the third was a Dell Latitude C600 852 MHz with 256 Mbytes of RAM. All three laptops used Orinoco Gold Wi-Fi cards by Lucent and ran MeshAP [1] mesh networking software that was bootable from a CD-ROM. One vehicle was a 2002 Volkswagen Passat, one vehicle was a 1995 BMW 318is, and one vehicle was a 1993 Saturn station wagon.

After rendezvousing at a roadside rest area, the three laptops were booted using the mesh networking software and subsequently configured in a mesh network. After making sure that each laptop could “see” each of the other laptops, all IP addresses were noted. Two console sessions were then initiated on each laptop, each session running “ping” commands to one of the other two laptops and writing the results of the ping commands to stored datafiles. The three vehicles then proceeded onto the highway and commenced driving in single file. There were very few large terrain features, with the road passing through gently rolling hills. Since the road was in interstate highway, there were no buildings or other large structures that could come between the test vehicles. The weather was completely sunny with no fog, rain, or clouds. Approximately every three miles, the rearmost vehicle sped up and passed the other two vehicles to take up a new position at the front of the caravan. No attempt was made to maintain the exact same distance between test vehicles. This process continued for about 30 kilometers until all vehicles exited the highway in single file, proceeded to another roadside rest stop, and the laptops were checked prior to their shutdown.

III. RESULTS

As this was primarily a “proof of principle” test, no attempt was made to establish detailed boundaries for network throughput, bandwidth, or node connectivity. For this first test, the use of the “ping” command was deemed adequate to establish the efficacy of constructing a mesh network for inter-vehicular communication and to determine a very general idea of the amount of time required to send data packets over a

vehicular wireless mesh network traveling at normal highway speeds.

Each line of data written by the ping commands to its corresponding datafile consisted of: 1) the number of octets sent over the network [64 octets for this test], 2) the IP address of the source node for the ping command, 3) a network “hop” counter that decremented once for each network hop, 4) a counter indicating how many individual ping commands had been sent, and 5) the time in milliseconds required for the 64 octets of ping command data to complete a network path from the sending node to the receiving node and back. Two files were written in each vehicle, one each for the network connection to each of the other two vehicles. A total of six files were written, although two were later found to be corrupted and unusable for analysis. The remaining four files contained data on the ping commands from each of two vehicles to the other two vehicles.

A. Datafile results

Each of the four columns of values from Table 1 (below) represents results from one datafile from one vehicle during the drive test. Therefore, the values show the results of the ping commands sent from one network node (vehicle) to another network node (vehicle). With three vehicles, there should have been six datafiles and therefore six columns of data, but the corrupted datafiles removed two of those possible six sets of values.

B. Network connectivity

Table 1 values show that the packet loss rate between nodes of the network for the four usable datafiles was: 3.61%, 4.97%, 6.66%, and 10.41%. This indicates that the wireless mesh network was connected throughout almost the entire drive time, although the packet loss for each pair of nodes was not very consistent. The difference between packet loss rates can be ascribed to the

differences in position of the laptop within each vehicle and the position of each vehicle with respect to the other vehicles.

C. Packet latency averages

Table 1 also shows the packet latency (as ping response time) for each of the four usable datafiles. The overall average ping response times for the four usable datafiles were: 2.4 ms, 2.5 ms, 2.5 ms, and 2.5 ms. These values were very consistent across all four usable datafiles.

There were times during three of the usable datafiles in which the ping response times rapidly jumped from the “normal” value of 2-5 milliseconds up to values between 600 ms to over 18350 ms, each of which also corresponded to times during which network packets were lost. These response times were not included in the calculation of the average ping response times shown in table 1.

D. Packet latency over time

Ping response time varied continuously over the complete drive time of approximately 15-20 minutes. Figure 1 (below) shows an example graph of the round-trip ping time from one node to another over the drive time, with drive time measured in ping sequence numbers (starting at 0 and ending at 1517). For this particular datafile, there were no instances of the ping round-trip transit time jumping from values of 2-5 milliseconds to thousands of milliseconds, as can be seen by all values lying within the range of 0-25 milliseconds on the graph.

Figure 1 shows that there is a minimum round-trip ping time value of about 1.8 milliseconds and a maximum value of slightly over 20 milliseconds. Figure 1 also shows that there are noticeable patterns in ping response time over periods of several seconds where the response time increases above the minimum value and stays there.

Datafile (arbitrarily labeled)	A	B	C	D
Packet loss rate (in percent)	3.61%	4.97%	6.66%	10.41%
Response time (in milliseconds)	2.4 ms	2.5 ms	2.5 ms	2.5 ms

Table 1. Network packet loss rate and response time




Figure 1. Ping Response Time vs. Drive Time

IV. CONCLUSIONS

The overall goal of the experiment was accomplished; a wireless mesh network for inter-vehicular communication was rapidly and inexpensively constructed and tested. The wireless network was generally stable, with only short periods in which inter-node connectivity was down.

Results of a simple analysis of the collected data show that, under generally normal highway driving conditions (standard speeds and distances) it was possible to maintain connectivity of the wireless mesh network. Also, the time required to send very basic network data packets (64 octets of a ping command) is relatively small and very consistent.

The results of the experiment show that it is relatively fast and easy to set up a wireless mesh network between vehicles traveling at highway speeds. The construction of such a network opens up many avenues by which to impact or measure vehicle flow and information transmission. For example, if all (or even merely the percentage) of vehicles traveling on a highway were known to have mesh networked connections, it would be a simple matter for a governmental agency to set up a roadside radio receiver that would instantly and continuously monitor the number of vehicles using that particular section of road. Another application relates to individuals who would like to use inter-

vehicular mesh networks within a small group of vehicles traveling together. A wireless mesh network would allow occupants of those vehicles to carry out such useful activities as streaming music from one vehicle to the others, or, using voice-over-IP technology, to speak to each other without incurring the bandwidth overhead or financial charges of cellphone use.

It was hoped that this experiment would begin to lay the foundation for the implementation of wireless mesh networks for inter-vehicular communication. While the data collected during the actual drive-time experiment was very rudimentary, the end results showed that such networks can be constructed very easily and can be appropriate for wirelessly connecting vehicles moving in a normal state on highways.

REFERENCES

- [1] LocustWorld home page, 2003. MeshAP software, <http://www.locustworld.com/>, referenced July, 2003.
- [2] Mitre Corporation, 2003. "Providing solutions for mobile adhoc networking", http://www.mitre.org/tech_transfer/mobilemesh/, referenced March, 2003.
- [3] Moore, J., Avnet, J., Ellis, A., 2002. "WikiWikiWan", <http://wiki.haven.sh/index.php/WikiWikiWan>, referenced March, 2003.
- [4] Verma S. and Beckman, P., 2002. "A Framework for Comparing Wireless Internet Service Providers with Neighborhood Area Networks", Proceedings of the Americas Conference on Information Systems 2002, Dallas, TX, August 8-13, 2002.