School of Engineering and Applied Science Real-Time and Embedded Systems Lab (mLAB)

University of Pennsylvania

Year 2005

A Multi-hop Mobile Networking Test-bed for Telematics

Rahul Mangharam* Jacob Meyers † Dr. Ragunathan Rajkumar ‡ Daniel D. Stancil** General Motors Corporation †† General Motors Corporation §

```
*University of Pennsylvania, rahulm@seas.upenn.edu
```

††

‡‡

This paper is posted at Scholarly Commons.

http://repository.upenn.edu/mlab_papers/36

[†]Carnegie Mellon University

 $^{^{\}ddagger}$ Carnegie Mellon University

 $^{{}^{**}{\}rm Carnegie}$ Mellon University

A Multi-hop Mobile Networking Test-bed for Telematics

Rahul Mangharam, Jacob J. Meyers, Dr. Ragunathan Rajkumar, Dr. Daniel D. Stancil

{rahulm, jjmeyers, raj, stancil}@ece.cmu.edu
General Motors Collaborative Research Laboratory
Department of Electrical & Computer Engineering
Carnegie Mellon University
Pittsburgh, PA. USA

Jayendra S. Parikh, Dr. Hariharan Krishnan, Christopher Kellum

{jayendra.s.parikh, hariharan.krishnan, christopher.kellum}@gm.com General Motors Corporation Warren, MI. USA

Copyright © 2004 SAE International

ABSTRACT

onboard vehicle-to-vehicle multi-hop wireless networking system has been developed to test the realworld performance of telematics applications. The system targets emergency and safety messaging, traffic updates, audio/video streaming and commercial announcements. The test-bed includes a Differential GPS receiver, an IEEE 802.11a radio card modified to emulate the DSRC standard, a 1xRTT cellular-data connection, an onboard computer and audio-visual equipment. Vehicles exchange data directly or via intermediate vehicles using a multi-hop routing protocol. The focus of the test-bed is to (a) evaluate the feasibility of high-speed inter-vehicular networking, (b) characterize 5.8GHz signal propagation within a dynamic mobile ad hoc environment, and (c) develop routing protocols for highly mobile networks. The test-bed has been deployed across five vehicles and tested over 400 miles on the road.

INTRODUCTION

The commoditization of high-speed wireless interfaces and low-cost Global Positioning System (GPS) devices provides the opportunity to deploy a range of useful and practical inter-vehicular communication applications. A test-bed has been developed primarily to evaluate the feasibility of high-speed networking between vehicles for emergency and safety notification and to support multimedia telematics applications. The focus is to develop and deploy a multi-hop wireless routing platform to facilitate on-road testing. We observe significant divergence in the performance of multi-hop wireless routing protocols between simulation and implementation and therefore stress the importance of on-road testing. In addition, channel measurements and propagation analysis have been carried out to further understand the nature of the communication environment. This paper discusses the platform deployment process and shares our driving experiences. Based on the network performance analysis

and insights from the test-bed, a vehicular networking architecture is being developed across multiple networking layers targeted specifically for inter-vehicular communication.

Consider the case when a vehicle accident on a freeway blocks two of the three available lanes. Most freeways rely on other drivers to notify the highway patrol, which in turn updates the freeway management system and messages are eventually broadcast on electronic signs along the highway. While the turnaround duration for accident notification may take considerable time, the traffic buildup is rapid and consequent blocking of emergency vehicles is still an unsolved and pressing problem.

Now consider the response when a subset of the vehicles has onboard inter-vehicular networking capability. From the accident site, alert messages are broadcasted and directed to all approaching vehicles. As the range of each wireless interface extends to less than a few hundred meters, it is essential to employ a multi-hop routing protocol to communicate 2-3 miles down the freeway. By using this mechanism to broadcast the event and its position information, approaching vehicles are notified of the hazard and possible congestion. As the broadcast message ripples through the traffic at high speed, approaching vehicles can use the next exit to plan a detour.

While traffic and alert notification systems such as OnStar [1] have been deployed, they are centrally managed entities. In the above example, we observe that the problem is local and requires action within the vicinity of the accident. Attempting a cellular call (used by OnStar) to each of the hundreds of vehicles in a particular region is not cost effective. Furthermore, the alert notifications are relevant only during the event lifetime and should be targeted specifically to approaching traffic.

In the following section, we discuss the range of vehicular applications and their unique networking challenges. We then describe the multi-hop mobile vehicular test-bed, our design decisions and driving experiences.

MULIT-HOP VEHICULAR NETWORKING

The primary goal of the multi-hop mobile test-bed is to provide insights towards the design of networking infrastructure necessary to support a range of intervehicular applications.

APPLICATION CATEGORIES

We focus on four application categories that encompass a broad range of possible services.

Emergency and Safety Messaging

Messages reporting vehicle accidents, sudden breaking, oil spills and other critical events need to be disseminated instantly and for the duration of the event's lifetime. As shown in Fig. 1, such messages are broadcast but within the scope of relevant vehicles approaching the event's position. Emergency messages are pertinent within the relative region of the event. This may be determined by restricting the message to be accepted and forwarded by only those nodes within a geographic region. In addition, navigation information such as vehicle direction and planned route may be leveraged. As the messages are critical, they must be delivered with the highest packet priority.

Traffic and Congestion Updates

These are semi-critical messages with updates on road conditions and congestion information. They may be delivered by fixed infrastructure-based transceivers or via mobile gateway vehicles equipped with both inter-vehicular wireless interfaces and cellular-data connections. Traffic and congestion updates are pertinent in the general scope of a region. A region may be described by a set of GPS coordinates. For example, a circular region may be described the GPS coordinate marking its center and the radius in meters. All vehicles within this region will accept and forward updates. Vehicles outside the designated region will drop any region-specific updates received. Updates will be delivered as broadcast packets.

Multimedia Telematics Applications

Feature-rich and interactive applications such as intervehicular voice communication, video streaming, file




Figure 1. Vehicle-to-vehicle multi-hop networking showing each vehicle's wireless range

transfer and collaborative driving applications fall within the scope of telematics applications. Using such applications, users can form or join public and private groups of vehicles. For example, a group of friends driving on a day trip will be continuously informed of the position of their friends' cars and can maintain a voice and chat channel to communicate collectively to the group. Such applications may be scaled to the enterprise level for fleet management or connect all vehicles heading to a particular destination. Telematics applications require a robust network and transport layer where connections are reliable and may be suspended and resumed smoothly based on the connectivity between interested parties. Furthermore, as all vehicles are connected across one or more hops, it is essential that the routes selected satisfy minimum admission control requirements and are self-healing.

Commercial Announcements

Travelers may choose to subscribe to commercial announcements regarding parking lots with empty slots, regional boarding and lodging information and other travel-related advertisements. Such applications require a subscription service and also a lightweight acknowledgement scheme to estimate the size of the targeted population.

In order to analyze the performance of the above application categories, we evaluated an application from each category over our test-bed. The next section describes the test-bed design followed by the unique challenges vehicular networks pose over traditional topology-based ad hoc networks.

TEST-BED DESIGN AND DEPLOYMENT

TEST-BED HARDWARE

The multi-hop wireless vehicular networking test-bed developed at Carnegie Mellon University (Fig. 2) employs a Differential GPS (DGPS) receiver with a magnetically mounted antenna. An onboard computer with a modified mini-PCI IEEE 802.11a wireless [2] interface forms the main entity of the setup. The physical layer has been modified to emulate the Dedicated Short Range Communication (DSRC) [3] standard specifications with a



Figure 2. Portable multi-hop mobile test-bed with DGPS receiver and onboard computer

10MHz signal bandwidth and operates at a variable carrier frequency which includes the 5.85 – 5.925 GHz spectrum.

Vehicles communicate with each other via a magnetic-mounted wireless antenna. The range of the wireless link is approx. 300m for line-of-sight reception. To facilitate multimedia applications, the test-bed includes a voice headset and a camera. All devices are powered by the vehicle's DC power system via the cigarette lighter, utilizing DC-DC power converters as needed. The equipment fits neatly in a plastic molded case and is easy to carry and quick to set up.

TEST-BED SOFTWARE

All onboard computers run RedHat 9 Linux (kernel 2.4.18-3) as this provides a fertile platform for network protocol and application development. We primarily use three layers of software on the test-kit. The software is built from open source libraries and is available for free.

Mapping and Communication Software

In order to visualize the current position of the vehicle, we adapted the RoadMap tool [4] for our test-bed. We added runtime display capability for multiple vehicles so the movement of each vehicle can be tracked as we drive. We added communication capability so that each vehicle's

onboard computer acts as a server and accepts connections from other vehicles. Each vehicle runs a User Datagram Protocol (UDP) client thread to connect to all other machines in the test-bed. As the connections are at the socket level, the application manages the end-to-end data exchange between the client and server. The underlying kernel-based networking software handles multi-hop routing along the set of links between the client and server. Using this client-server setup, each vehicle exchanges its GPS information (position, speed, direction, etc.) and its network information (packet sequence number, fragmentation, frame length, etc). implemented our own packet headers for efficient exchange of position and network information. GPS coordinates are computed five times every second and have an accuracy of ≤ 2m. Consequently packets are exchanged five times every second.

All GPS and network information exchanged between vehicles is logged by each machine. This enables us to playback the route driven and visualize the vehicles on a vector-based rendering of the map traversed. The maps use TIGER/Line 2002 data files available for free from the U.S. Census Bureau [5]. In Figure 3, we observe the playback of a trip with five vehicles (top right) illustrating the vehicles connected, a panel to send emergency messages, a playback control (bottom left) to speedup or




Figure 3. RoadGPS vehicle visualization with GPS tracking, network connectivity and audio/video/messaging communication tool. A playback control is provided for viewing logged trips.

slow down the playback, client and server connections and a multimedia application with voice and video (top left)

Table I lists the current functionality of our mapping and communication application – RoadGPS. All position and networking information within the multi-hop network on the road can also be channeled to the Internet via the 1xRTT cellular connection. This way, our team at the GM Technical Center in Warren, MI can monitor all of the vehicles and the data transmitted between them as they drive in Pittsburgh, PA or anywhere else with cellular network access. The ability to monitor the network in real-time assures that data is correctly being logged and allows real-time network performance and signal propagation analysis.

Multi-hop Networking Software

We tested several existing ad hoc routing protocols such as DSR [6] and AODV [7] kernel implementations. While the protocols provide connectivity across multiple hops, their performance in a highly mobile environment was unreliable. Most ad hoc routing protocols have been tested primarily through simulation or in small test-beds (< 10 nodes) with low mobility. In a highly mobile environment such as inter-vehicular communications where connectivity changes often, these protocols were unable to reconstruct routes fast enough.

To improve performance, we modified the routing protocol to select paths based on link stability rather than hop-count to dampen the rapid link oscillation when vehicles are in contact only briefly. Several conceptual changes are required to tailor multi-hop routing for vehicular networks and they are discussed in the following sections.

Multimedia Conferencing and Application Software

Audio and video streaming was carried out using conferencing software based on the H.323 standard libraries. The reception of the audio and video codec was clear but suffered from an extended delay when reconnecting temporarily broken links. New buffering implementations to address the frequent but brief link connection are necessary. In addition, a new session layer protocol that can suspend an open connection when the link is briefly disconnected and resume the flow upon subsequent connection would be very useful.

Table IFunctionality of RoadGPS Vehicular Networking Tool

1	Map & Display multiple vehicles
2	Communicate with multiple vehicles
3	Communicate over multiple hops
4	Send Safety Messages
5	Stream data & music files
6	Communicate using Cellular 1xRTT
7	Complete trip logging and playback
8	Analysis & Graphing functions

DRIVING EXPERIENCES & ANALYSIS

The multi-hop networking test-bed has been deployed across five vehicles provided by GM. The group of vehicles has been driven over 400 miles each. We chose four environments: urban (densely populated and crowded with several high-rise buildings such as downtown Pittsburgh), rural (flat with fairly open roads), highway (high speed driving along I-79) and city driving through Pittsburgh. In general we were able to maintain connectivity across all vehicles over 1.2km using multi-hop routing. This shows the basic usefulness of multi-hop networking as the range of each vehicle's wireless connection is limited to 300m under good conditions. The terrain of the city of Pittsburgh is quite hilly and offered us several non-line-of-sight opportunities where multi-hop routing proved to be very useful in maintaining connectivity.

We have developed a data analysis toolkit in MATLAB to analyze connectivity, signal-to-noise ratio, error rates, and data rates over different distances, speeds and link transmission rates. A comparison of logged packets transmitted from each local node with packets received by each of the other nodes in the network, combined with the GPS data from both nodes, allows a detailed statistical analysis of the network's connectivity performance with regard to dropped packets. The network connectivity performance for each node, as measured by packet reception, can be analyzed as a function of absolute and relative node speed, the distance between nodes, transmission signal strength, and data transmission rate. For example, Fig. 4 shows that the percentage of total sent which are dropped increases exponentially as the distance between nodes increases for data transmitted at a data rate of 6MBPS, as one might expect.

The 5.8 GHz signal propagation characteristics within the dynamic mobile ad hoc environment can be determined




Figure 4. Percentage of total packets sent which are dropped vs. distance




Figure 5. Total Average Measured Signal Attenuation and Free Space Propagation Model vs. Distance

using calibrated measurements of received and transmitted signal strengths in conjunction with the GPS data. The signal attenuation between each node can also be determined and modeled as a function of absolute and relative node speed, the distance between nodes, transmission signal strength, and data transmission rate. For example, the initial data shown in Fig. 5 suggests that signal attenuation can be roughly modeled using the free space propagation model

$$P_s \cong \frac{1}{r^n} + Offset$$

where r = distance between nodes, and n \approx 2, as one might expect.

Fig. 6 provides a graphical representation of the routing table over time for source node with IP address 192.168.1.4 (referred to as IP4 here) attempting to communicate with destination node IP2 across one or more intermediate nodes. The y-axis plots the connectivity to nodes with IP addresses IP1, IP2, IP3 and IP5 with respect to IP4. A connection to IP4 is represented by a line passing through 4 on the y-axis. We observe that initially, IP4 is connected only to IP5 and sends the packet to IP5. IP5 forwards the packet to IP3 which in turn forwards it to IP2. This is the situation when the cars were parked in the order of their IP addresses and were in range of only their immediate neighbors. However, at a later time, we see considerable activity as cars overtake each other, or are separated by traffic lights or when direct connection is obstructed by passing large transportation carriers.

Our results are based on a modified implementation of AODV routing protocol. In the following section we outline several aspects where a new class of routing protocols needs to be developed to address the different aspects of vehicular networks.

UNIQUE CHALLENGES FOR VEHICULAR NETWORKS

Several designs for infrastructure-less "ad hoc" wireless networks have been proposed over the past decade [6].




Figure 6. Graphical representation of a vehicle's routing table over time

They predominantly belong to the class of Topology-based protocols where a source node, s, requires a communication path with destination, d. As the range of the wireless interfaces is limited, the communication path will generally traverse multiple intermediate nodes, each of which relay messages. Such protocols assume every node is described by a logical (named) address (e.g IP address) and require a full search of the network to find the destination node. This search procedure (or link flooding) generally does not leverage any topological structure of the network and therefore floods the network with broadcast requests to find the destination node. Examples of such protocols are: Dynamic Source Routing (DSR), Ad hoc On-Demand Distance Vector Routing (AODV) and Optimized Link State Routing Protocol (OLSR) [6].

Are Topology-based Routing Protocols suitable for Vehicular Networks?

Several Topology-based ad hoc routing protocols have been shown to work for small and lightly mobile networks. We outline several factors in the routing protocol requirements for vehicular protocols that are not addressed by existing protocols as the challenges and applications addressed are largely different.

1. High Mobility

Ad hoc networking protocols have traditionally been simulated with average speeds less than 30kmph [6] and have shown to perform poorly under higher mobility conditions. In contrast, a vehicular network is by definition always mobile with automobiles moving in excess of 140kmph and reaching relative speeds over 300kmph. This 10x difference of node speed requires rewriting the protocol to not be limited by (a) handshaking (e.g. routerequest/route-reply), (b) formation of end-to-end connection oriented communication along intermediate nodes (e.g. source routing), and (c) rebroadcasts of end-to-end route information (e.g. routebroken alerts). Furthermore, there is a need to test the correctness and performance of routing protocols through implementation rather than simulation. implementation and testing bring about anomalies such as

obstructions (buildings), highly dense networks (city driving), sparse networks (highway driving), etc.

2. Broadcast over Unicast

In a vehicular network, emergency messages and traffic updates are relevant to all nodes in the vicinity of the event. There is a need to communicate with nodes based on their **relative position**, destination, speed, direction, etc. The addressing scheme of the routing protocol must therefore take into account packet broadcast as the common way to communicate than using unicast messaging.

3. Geocasting with Topology Information:

As vehicle paths are constrained by the network of roads, connectivity and flooding will benefit from leveraging knowledge of the **traffic topology**. Most Topology-based protocols do not assume any known network structure and flood the entire network. This severely limits their scalability to less than 30 nodes. Furthermore, as messages have a geographic relevance, it is essential to use node position as a key network construct.

4. Intersection-Intersection Routing Vs. Vehicle-to-Vehicle Routing:

In order to exploit geographic information, several greedy location-aware routing protocols have been proposed. These however fall short when a road bends and the path to the destination is not as the crow flies. Greedy locationaware routing protocols attempt to reach an intermediate node only if it is geographically closer to the destination, and therefore the packet delivery rate is severely limited by their use of local information. On the other hand, using limited navigation information, packets can be routed from intersection-to-intersection rather than from vehicle-tovehicle. This eliminates the need to keep track of the addresses of intermediate hops or to maintain a destination-based routing table. Topology-based protocols do not exploit any such information and are based solely on logical addressing. Using intersection-based routing, vehicles are used opportunistically if they lie along the region of interest. Furthermore, by fixing routing to static intermediate points, vehicles moving outside the area of interest will not adversely affect the path.

5. Message Lifetime:

Topology-based protocols aim to eliminate duplicate messaging and stop sending a message once an acknowledgement is received. In a vehicular network, if an incident occurs at a particular location, there is a constant need to update all vehicles approaching the incident site. It is therefore necessary to continue broadcasting the message to all new nodes until the incident lifetime has expired. Furthermore, as the network connectivity may be sparse, it is useful to hold the message until a new neighbor is detected. Topology-based protocols employ a "hot-potato" approach to routing where once a packet is forwarded it is forgotten. Message lifetime support is crucial to maintaining alert status and also increasing message delivery rate in sparsely connected networks.

6. Dirty Routing Tables:

With vehicles traveling over 140kmph, a node 1km away will be within 5m of standing traffic in less than 25 seconds. Such quick changes in the network neighborhood will result in drastic changes in routing table information. For example, a vehicle leaving downtown and entering a highway will corrupt the routing tables of all highway nodes in its vicinity with information of nodes just traveling in downtown. There is a strong need to maintain location-based routing information and ensure its freshness. Topology-based routing protocols assume routing updates occur at a much slower pace.

ADDRESSING FOR VEHICULAR NETWORKS

We now focus on the addressing requirements for vehicular networks. In addition to sending messages targeted to a particular vehicle, there are several instances where a message is targeted to a group of vehicles that meet a particular criterion such as current position within region of interest, speed, direction, etc. Three addressing schemes have been identified to support named, geographic and node property-based addressing. The following addressing schemes will be incrementally deployed in the test-bed.

A. Named/Assigned Addresses:

For unicast messaging and to identify vehicle-type groups, Assigned Addresses provide a single vehicle logical naming such as IP address, VIN, MAC address, etc. An example is: "Message for Buick FJF2323"

B. Geographic Addresses:

For messages focused on a particular geographic region, vehicles may choose to accept, forward or drop packets based on their geographic attribute.

1. Absolute Geography

- **Direction**: "Message for All neighbors driving east on I-90"
- Navigation: "Message for All cars headed for Exit 22"
- Region: "Message for Downtown Pittsburgh"

2. Relative Geography

- Hop count: "Message for All cars within 6 hops radius"
- **Speed:** "Message for All cars moving within +/-8mph of my speed"

C. Property-based Addresses:

For messages focused on an attribute of a vehicle:

- 1. Connectivity: "Message for any mobile gateway"
- **2. Vehicle Type:** "Message for Trucks only, or GM vehicles only, or 2-wheelers
- **3. Vehicle Dynamics:** e.g. Absolute Speed: "Message for vehicles moving at 40kmph +/-5kmph"

We observe that the vehicular network routing protocol will need to support all three addressing schemes and ensure their coexistence. Therefore, based on the application, an appropriate addressing scheme will utilize broadcast, multicast groups and unicast for a well-defined subset of vehicles.

CONCLUSION

We describe the design and deployment of an experimental multi-hop mobile test-bed for vehicular networks. The goal of the test-bed is to get insights into the vagaries of the wireless channel and routing protocols to deliver both mission-critical emergency messages and interactive telematics applications. Our current deployment includes a DGPS receiver, an onboard computer, a modified 802.11a radio card, a 1xRTT cellular-data connection and multimedia peripherals. Our experiences from driving the test-bed over 400 miles indicate that traditional applications and network infrastructure are usable. However, to meet the demands for a scalable and stable vehicular network both routing protocols and transport protocols need to be developed to address the dynamism of a highly mobile and lightly connected network of vehicles. We provide a list of technical requirements for the design of routing protocols and highlight the unique addressing capabilities necessary. As future work, we aim to develop and deploy the ideas presented in this paper.

ACKNOWLEDGMENTS

The authors would like to thank Daniel Weller and Suchit Mishra for their valuable inputs.

REFERENCES

- 1. OnStar Safety System. http://www.onstar.com
- 2. IEEE 802.11a, Part 11: High-speed Physical Layer in the 5 GHz Band, Supplement to IEEE 802.11 Standard, September 1999.
- 3. Dedicated Short Range Communications. http://www.leearmstrong.com/DSRC
- 4. Roadmap. http://sourceforge.net/projects/roadmap/
- Census 2000 TIGER/Line Data. http://www.esri.com/data/download/census2000 tigerline/
- 6. J. Broch, et al. A performance comparison of multi-hop wireless ad hoc network routing protocols. In Proc. Mobicom, pages 85--97, 1998
- 7. Charles E. Perkins, Elizabeth M. Belding-Royer, and Samir Das. "Ad Hoc On Demand Distance Vector (AODV) Routing." IETF RFC 3561