UC Irvine

Title
Design, implementation, and test of a wireless peer-to-peer network for roadway incident exchange

Permalink
https://escholarship.org/uc/item/0f67v6ph

Journal
Proceedings of the 9th International Conference on Applications of Advanced Technology in Transportation, 9

Authors
Harmon, Trevor
Marca, James E.
Klefstad, Raymond
et al.

Publication Date
2006-08-01

Peer reviewed

Design, implementation, and test of a wireless peer-to-peer network for roadway incident exchange[1]

Trevor Harmon [2]   James Marca[3]   Ray Klefstad [2]   Peter Martini [2]

ABSTRACT

Vehicular traffic monitoring and control has had a strong infrastructure bias—data is collected centrally, processed, and then redistributed to travelers and other clients. There are several efforts to decentralize traffic monitoring by leveraging advanced local area wireless technology. This paper describes our implementation of such a traveler-centric system, called Autonet. Each Autonet client exchanges network knowledge wirelessly with other, nearby clients. It is demonstrated that knowledge about traffic state can be propagated using this system. The client programs were also used to test the actual throughput possible for messages sent from one vehicle to another using 802.11b wireless hardware. These measurements establish the maximum throughput at about 4,000 incidents for two vehicles moving in opposite directions at highway speeds.

INTRODUCTION

For the next generation of intelligent transportation systems, we envision a decentralized, autonomous communications network between vehicles that leverages the unlicensed spectrum provided by the government and the off-the-shelf hardware that has been developed to exploit this spectrum (802.11a/b/g, DSRC[4] band, etc.). The system we propose expects individuals to share their own travel data freely and allows each traveler to buy as much or as little a device as they need. We call this concept the *Autonet*.

This paper documents an initial implementation of the core vehicle-to-vehicle client application for Autonet. The primary goal is to move from simulation to the real world. The results document the *current* capabilities of hardware for vehicle-to-vehicle communication of traveler information, and they provide working protocols and communications parameters for future simulation studies. The in-vehicle client application will also be useful in future studies of driver interaction and market acceptance.

BACKGROUND

Each Autonet peer does two things: it identifies changes from baseline conditions, and communicates these changes to other vehicles. Vehicles must identify situations that are unexpected, without reference to a central authority, and must
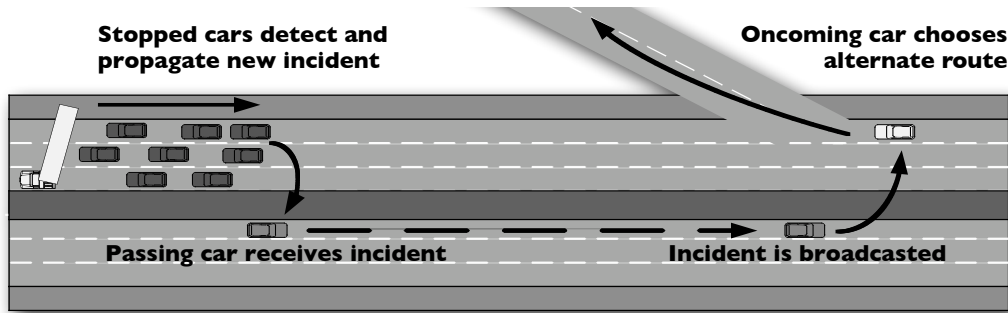
---

[2]University of California, Irvine, Department of Electrical Engineering and Computer Science. Irvine, California 92697.

[3]University of California, Irvine, Institute of Transportation Studies. Irvine, California 92697.

[4]Dedicated Short Range Communications, standardized by the U.S. Dept. of Transportation.

Figure 1. By running variations of this basic scenario on actual roadways, we could judge the prototype's effectiveness.

transmit this information to other vehicles. The Autonet protocols and systems arbitrate between the information gathered by all vehicles so that a correct interpretation of current conditions can be established and propagated.

We explored various ideas for communicating information in the Autonet and decided on the scenario shown in Figure 1. In this scenario, a message about traffic state is passed from one direction of traffic flow to the other for rebroadcast to upstream vehicles. We tested this scenario with a proof-of-concept prototype built entirely from COTS (commercial off-the-shelf) hardware. Although 802.11b seemed like the obvious choice for this prototype's wireless technology, this decision brought with it a host of questions and somewhat contradictory answers.

Ebner (Ebner et al. 2003) recommended against using 802.11b, stating that simulation results indicated that its performance would be unacceptable in urban environments and at high speeds. In contrast, Singh (Singh et al. 2002) conducted several field tests and concluded that 802.11b performance is "suitable for inter-vehicle communications." These papers are not in conflict, but rather interpret "suitability" in different ways. Ebner was striving for a seamless approximation of the Internet on the road over 802.11b, while Singh had a much more modest goal of vehicle-to-vehicle communication.

In other related work, Nadeem (Nadeem et al. 2004) designed a system around 802.11b and reported some results mixing real world tests and simulation for vehicles moving in the traffic stream. Wu (Wu et al. 2004) developed a detailed protocol for sending a message from one geographic area to another, which was tested in simulation using the 2 Mbps 802.11 DCF MAC protocol. Finally, Aziz (Aziz 2003) used 802.11b in infrastructure mode to enable vehicle-to-vehicle communication between vehicles with fixed IP addresses while traveling along a corridor.

The Autonet has modest requirements for vehicle-to-vehicle communication. It does not require a complete, multi-hop network to obtain quantitative improvements in the level of system awareness for participating drivers. Given that a single-hop network is all that is requied for the initial Autonet vision, we concluded that 802.11b was more than adequate for building our field test units.

DESIGN

As a first step in reaching the goals of Autonet, we have developed a prototype
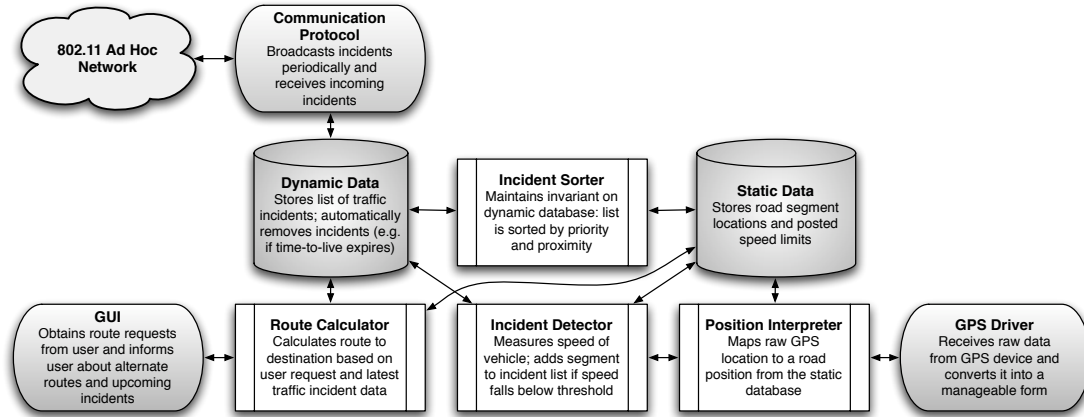
Figure 2.   The major software components of the Autonet prototype.

Autonet computer. This prototype demonstrates the feasibility of Autonet and provides a platform for testing our algorithms in a realistic environment on public roads using COTS hardware.

Software components

Figure 2 shows the major software components and the relationships between them. Each of these components is designed to operate concurrently with the others, performing operations on incidents as they come in and responding to user requests asynchronously. For simplicity, the components can be thought of as operating in a sequential fashion as in the following scenario:

1. The Communication Protocol receives a message from a passing Autonet vehicle. The message describes the location, time observed, and severity of the incident.
2. The Dynamic Data component receives the new incident and checks it for validity. Valid incidents are placed into dynamic storage.
3. The Incident Sorter monitors the Dynamic Data for changes and re-sorts it according to some evaluation function. This function takes into account the spatial and temporal proximity of the incident, and its severity, among other criteria.
4. The Route Calculator monitors the Dynamic Data for changes and recalculates the shortest route to the current destination, incorporating new incident information.
5. The Communication Protocol queries the database for the highest priority incidents, as this set may have changed with the addition of new incidents. It broadcasts the set to other Autonet-enabled vehicles.

Communication protocol

The communication protocol impacts the overall quality of the system. It is responsible for exchanging incidents with other Autonet-enabled vehicles, and it ensures that users have accurate knowledge of the true traffic state. It accomplishes these tasks by broadcasting a list of incidents—segments of road observed to be slower than expected—while simultaneously listening for incident reports from the surrounding vehicles.

The message format for incident exchange is simple; the decision of which incidents to send is not. The problem is that over time, a group of Autonet-enabled vehicles may build up a list of tens of thousands of incidents. According to our measurements (see Section 6), two vehicles passing each other in opposite directions can each exchange only a few thousand incident messages with 802.11b technology, even under the best of conditions. Therefore, the communication protocol must decide which portion of the incident list to broadcast. We address this problem by assigning a priority to each incident, calculated with an evaluation function that takes into account:

- **Deviation from expected speed** The higher the deviation from the expected speed, the more important is the incident. Currently, we define expected speed as the posted speed limit for the road, but for future prototypes we will define it more accurately as the historical average of vehicles passing on the road segment.
- **Size of the road** Highways have higher priority than local roads, etc.
- **Distance from local position** Nearby incidents are more likely to have a direct impact on travel than distant incidents.
- **Time** The older the information, the less accurately it describes current conditions.

The Incident Sorter applies its evaluation function to sort the incident data by priority. The communication protocol then cycles through the highest priority incidents and broadcasts them continuously for passing vehicles to collect.

IMPLEMENTATION

To show that the Autonet concept is feasible, we constructed prototypes of an Autonet device and installed them in three vehicles. These prototypes helped us perform experiments in a real-world environment using COTS hardware: basic laptop computers, an 802.11b wireless Ethernet card in ad hoc mode, and GPS receivers that communicated with the laptops over Bluetooth. For software, we implemented the components shown in Figure 2 in the Java language.

Incident detection

Our Incident Detector, as discussed in Section 4, provides a basic algorithm for detecting new incidents. The inputs are speed limit information from the static database, and the current speed obtained from the Position Interpreter. An incident is declared if the vehicle drops below 80% of its expected speed. Although this technique works well for freeway driving, it is inadequate for roads with flow interruptions such as intersections and traffic controls.

Although we could have attempted to mitigate these problems, much work has already been done in the transportation research community on incident detection, and we did not wish to re-invent the wheel for our initial prototype. As we discover a known set of capabilities and resources for an Autonet peer, we will then turn to integrating the rich body of literature on incident detection techniques into the unique requirements of the Autonet devices.

User interface

To the casual viewer, the display of an Autonet device would appear no different from commercial navigation systems that are becoming more popular. The
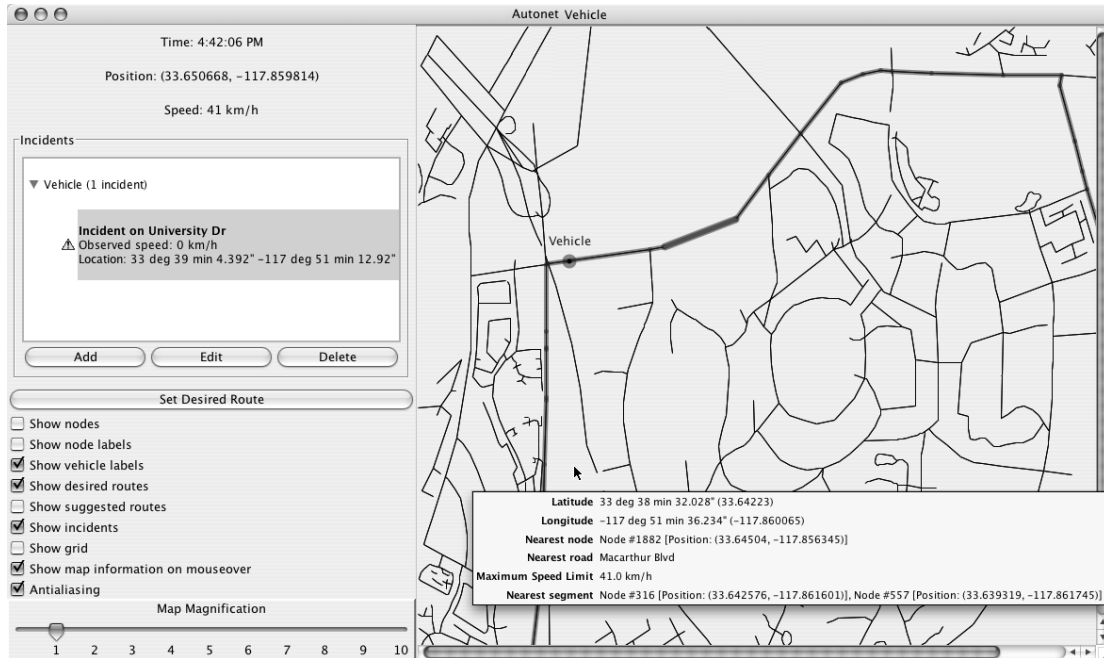
Figure 3. The prototype shown in this screenshot is used to explore new feature ideas and to verify the performance and correctness of Autonet.

Autonet user interface would improve upon these more conventional systems, however, by displaying up-to-the-minute traffic information collected from other Autonet-enabled vehicles. By combining this information with the driver's chosen route, Autonet's interface can then display a more appropriate route (if found) superimposed over a map of the vehicle's surrounding area.

Figure 3 shows our prototype's user interface. The left side of the screen shows the vehicle's speed and position as reported by the GPS device, and a list of known incidents. This list includes directly observed incidents, and second-hand information received from other Autonet vehicles. At the bottom are controls for adding, editing, and removing incidents and for manipulating various attributes of the map display. The right side of the screen shows the road map of the vehicle's location. The map shows the vehicle's position and its chosen route.

EMPIRICAL RESULTS

To benchmark the prototype, we focused on the real-life scenario shown earlier in Figure 1. Three vehicles (at minimum) needed to be equipped with our Autonet prototype to test this scenario. We successfully executed a number of variations of this scenario on streets and freeways in Irvine. In addition, we ran a variety of custom benchmarking programs for testing the performance of 802.11b technology in the Autonet environment. Figure 4 shows the result of this effort.

For Figure 4(a), we tested the ability of 802.11b to transfer incidents between two vehicles that pass each other in opposite directions at identical speeds. The horizontal axis refers to the *absolute* speed of the vehicles. The plot shows the average number of incidents that each vehicle could send over four trial runs.
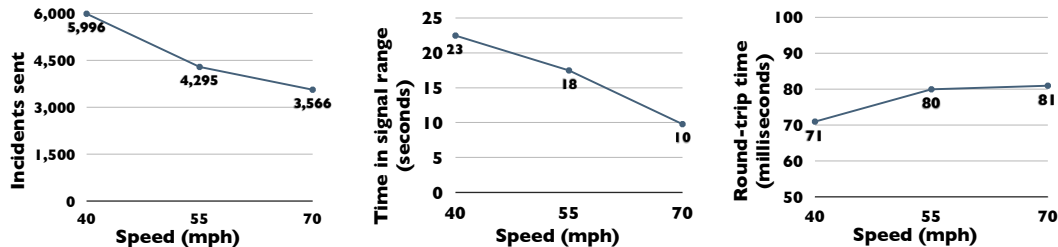
(a) Raw throughput for incident transfer.

(b) Time the vehicles remained within signal range.

(c) Round-trip time for query/response.

Figure 4. These graphs show our benchmark test results of 802.11b communication, aided by external antennas, between vehicles moving at identical speed in opposite directions on public roads.

As expected, the number of incidents exchanged drops sharply as the speed of the vehicles increases. The decline is largely due to the shorter time in which the two vehicles are within signal range, but we suspect that the Doppler effect also plays a role in limiting throughput with increasing speed. We note, however, that even when two cars travel at highway speed—70 miles per hour—they are able to send and receive over 3,000 incidents each. These are very promising results, as they show that one vehicle can inform another about every single segment of road for a typical urban area of 64 square kilometers. If we assume that one out of every five segments is an incident, the area grows to 324 square kilometers.

For Figure 4(b), we measured the amount of time during which two vehicles, when traveling in opposite directions, could remain in 802.11b signal range of each other and successfully exchange incidents. The surprisingly long periods during which the two vehicles were able to remain in contact can be attributed to the external antenna mounted on each vehicle's roof. Although we had hoped to collect data without this antenna for comparison purposes, we were unable to exchange incidents reliably after removing it, even at slow speeds. The consequence for Autonet is that any production version of the device would require the consumer to install an external antenna atop their vehicle.

Finally, Figure 4(c) shows the round-trip time for message exchange using 802.11b. Here, we define "round-trip time" as the time required for one vehicle to send an incident request to another and fully receive a response to the query. Based on these measurements, combined with the measurements from Figure 4(b), we calculate that when using a request/response method of incident exchange, a minimum of approximately 100 incidents can be received per vehicle. (That is, at the worst case of 70 mph, the vehicles are in signal range for ten seconds, so 10 seconds ÷ 81 milliseconds = 123.)

CONCLUSION

This paper presents the results of a real-world implementation of a vehicle-to-vehicle communication system which enables decentralized traffic information propagation. Future work will focus on defining the minimum requirements and capabilities of an Autonet peer. These capabilities will be developed in tandem

with simulation efforts aimed at determining more robust communication protocols and incident detection algorithms.

REFERENCES

Aziz, F. M. (2003). "Implementation and analysis of wireless local area networks for high-mobility telematics." Master's thesis, Virginia Polytechnic Institute and State University.

Ebner, A., Rohling, H., Wischhof, L., Halfmann, R., and Lott, M. (2003). "Performance of UTRA TDD Ad-Hoc and IEEE 802.11b in vehicular environments." *57th IEEE Semiannual Vehicular Technology Conference*, Vol. 2, IEEE, Jeju, South Korea. 960–964.

Nadeem, T., Dashtinezhad, S., Liao, C., and Iftode, L. (2004). "TrafficView: Traffic data dissemination using car-to-car communication." *Mobile Computing and Communications Review*, 8(3), 6–19.

Singh, J., Bambos, N., Srinivasan, B., and Clawin, D. (2002). "Wireless LAN performance under varied stress conditions in vehicular traffic scenarios." *Proceedings. IEEE 56th Vehicular Technology Conference, Fall 2002*, Vol. 2. IEEE, 743–747.

Wu, H., Fujimoto, R., Guensler, R., and Hunter, M. (2004). "MDDV: a mobility-centric data dissemination algorithm for vehicular networks." *Proceedings of the first ACM workshop on Vehicular ad hoc networks*, New York, NY. ACM Press, 47–56.