

TrafficView: Towards a Scalable Traffic Monitoring System

Tamer Nadeem, Sasan Dashtinezhad, Chunyuan Liao, Liviu Iftode
Department of Computer Science
University of Maryland
{nadeem,sasan,liamay, liviu}@cs.umd.edu

Abstract

Vehicles are part of people's life in modern society, into which more and more high-tech devices are integrated, and a common platform for inter-vehicle communication is necessary to realize an intelligent transportation system supporting safe driving, dynamic route scheduling, emergency message dissemination, and traffic condition monitoring. TrafficView, which is a part of the e-Road project, defines a framework to disseminate and gather information about the vehicles on the road. Using such a system will provide a vehicle driver with road traffic information, which helps driving in situations as foggy weather, or finding an optimal route in a trip several miles long. This paper describes the basic design of TrafficView and different algorithms used in the system.

1 Introduction

Vehicles are part of people's life in modern society, into which more and more high-tech devices are integrated. Most of the current research focuses on the functionalities of individual vehicles, and less attention has been paid to the cooperation among vehicles and road facilities, which forms the whole transportation system. Moreover, a common platform for inter-vehicle communication is necessary to realize an intelligent transportation system supporting safe driving, dynamic route scheduling, emergency message dissemination, traffic condition monitoring, and etc.

The *e-Road* project is an attempt to achieve the aforementioned goals by providing a scalable infrastructure for inter-vehicle communication. Specifically, the e-Road project is aimed at building a system with the following characteristics:

- *Real-time message dissemination platform:* The messages can be efficiently delivered from moving vehicles or fixed stations to other vehicles, with small delay and low bandwidth cost. The dissemination system could be used in sending messages about traffic condition monitoring, road condition, accident report, road-side e-advertisements, etc.
- *Information query platform:* Besides "passively" receiving messages, a vehicle is enabled to query information about specific targets. For example, a vehicle

can query about the average vehicle speed in a region, road condition at Exit 11, or hotels 10 miles away.

- *Reliable information exchange protocol:* This is a necessary set of algorithms and metrics for connection-oriented applications. For example, it can be used for music downloading, back-seat passenger games, or connection to the Internet.

In this paper, we present *TrafficView*, which is a part of the e-Road project. TrafficView defines a framework to disseminate and gather information about the vehicles on the road. Using such a system, a vehicle driver will be aware of the road traffic, which helps driving in situations like foggy weather or finding an optimal route in a trip several miles long.

A GPS receiver shows a static view of the map, whereas TrafficView provides the driver with a dynamic view of the road traffic, and therefore complements the GPS receiver. When integrated with the traditional digital map system, TrafficView would be able to provide the functionality of real-time automatic route scheduling. Moreover, in such a platform, other applications such as accident alert, and road-side e-advertisement can be easily implemented.

This paper describes our experience in developing the TrafficView system. Throughout our experimentation, we performed a detailed study of different information dissemination techniques under various road density and vehicle mobility conditions.

The rest of the paper is organized as follows: The next section presents the background, and the description of the problem is given in Section 3. In Section 4 and Section 5 we describe the design of TrafficView and the mechanisms used in the system. The System performance is studied in Section 6. In Section 7 we summarize the related work. Finally we present our conclusions and future work in Section 8.

2 Background

2.1 Mobile Ad-hoc Networks

Pervasive computing, a hot topic in recent systems research, gives researchers a new broad horizon to explore. At the heart of the pervasive computing, distributed system and mobile computing have received a lot of attention [15]. With the development of the pervasive computing, the use of computers will become invisible and natural, and be tightly combined with physical objects to form a *data space*, where the data is inherently dispersed and connected as a part of physical objects.

Wireless Mobile Ad Hoc Networks (MANETs) are one of the contributions toward the aforementioned goal; quite attractive because of higher flexibility and ad hoc infrastructure. However, they also give rise to challenges to researchers, because of the inherently-highly-dynamic network topology and limited transmission range in these networks. A lot of work has been done, involving routing, data caching, data querying, and so forth related to MANETs.

2.2 Inter-Vehicle Communication

The research in Inter-Vehicle-Communication has emerged in the past couple of years; mainly because it is a good experiment platform for Wireless Mobile Ad Hoc Networks, and has a great market potential [8]. Several major automobile manufactures and universities have begun to investigate in this field; GM research center in CMU [7], BMW Research Labs [16] and Ford Research Labs [11], Rice University [17][13], and Harvard University [4] are a few to name.

In addition to the similarities to MANTET such as short radio transmission range, low bandwidth, omnidirectional broadcast (at most times) and low storage capacity, inter-vehicle communication has its unique characteristics and challenges as well:

- *Rapid changes in link topology.* Because of the relative movement of the vehicles, the connectivity between vehicles is always changing. For example, if vehicles' speed is around 60 mph, i.e., 25m/s, and the transmission range of wireless networks is 250m, the connectivity between two vehicles could last for at most $500/25 = 20$ seconds.
- *Frequently disconnected network.* In the case of low vehicle density, the gap between two vehicles might be several miles, far beyond the transmission range of wireless networks. In turn, the disconnection time could be minutes. Due to the fast movement of vehicles, and high dynamic traffic conditions, this situation is not uncommon.
- *Data compression/aggregation.* Wireless networks have a limited available bandwidth. In order to build a scalable system, data compression/aggregation mechanisms are required to save the bandwidth.
- *Prediction of vehicle's positions.* Vehicles normally run along pre-built roads, which remain unchanged over years. Therefore, given the average speed, current position, and road trajectory of a specific vehicle, the future position of that vehicle can be predicted.
- *Energy is not a big issue.* In sensor networks, the nodes are battery-powered, and it is not easy to replace the battery after deployment. This limits the running time of a node in such a network. Therefore, a lot of effort has been made to conserve energy in sensor networks. On the other hand, in a vehicle network, the vehicle itself can be used as a source of electric power, and therefore, energy is not a big issue in such a network.

2.3 Data Delivery Modes

In multihop wireless networks, such as vehicle network, data could be delivered using two different modes: *push* mode and *pull* mode.

In push mode, which is the focus of this paper, data is broadcast and propagated to all possible recipients that passively listen to the channel. There is no notion of route establishment or interaction between the senders and the receivers. Data is transported in unreliable and out of order fashion. Typical application scenarios are




Figure 1: The problem this paper addresses (a) and the diffusion mechanism (b and c)

vehicle accident alert, peering vehicles monitoring, road-side e-advertisements, and so forth.

In contrast, in pull mode the data is transmitted between a *query sender* (QS) and a *query responder* (QR), and there is some interaction between the two. The QS needs to actively send out a request to a specific target—a vehicle or a group of vehicles in a region—which in return reply with the available answer. To establish this session, the QS and QR should identify each other with their geographic position, unique identification number or other means.

3 Problem Description

Given a set of moving vehicles on the road, the goal is to exchange information about the position and speed of those vehicles among them to enable each individual vehicle to view and assess traffic and road conditions in front of it. As the vehicles move along the road, they might enter the transmission range of some vehicles, and exit that of others. Figure 1 (a) shows an example of a road with four lanes, on which four vehicles are moving. Two main mechanisms could be used to achieve this goal: *flooding* and *diffusion*. In the flooding mechanism, each individual vehicle periodically broadcasts (pushes) information about itself. Whenever a vehicle receives a broadcast message, it stores it and *immediately* forwards it by rebroadcasting the message. Obviously, this method is not scalable, due to messages flooding over the network, especially in high density roads.

In the other mechanism—the diffusion mechanism—each vehicle broadcasts information about itself and the other vehicles it knows about. Whenever a vehicle receives broadcast information, it updates its stored information and defers forwarding the information to the next broadcast period, at which time it broadcasts its updated information. The diffusion mechanism is scalable, since the number of broadcast messages is limited and they do not flood the network. We use the diffusion mechanism in TrafficView.

As an illustration of the diffusion mechanism, assume that in Figure 1 (a), vehicles

2 and 3 are in the transmission range of vehicle 1. Likewise, vehicles 3 and 4 are in the radio range of vehicles 2 and 3, respectively. At the beginning, each vehicle knows only its own position and speed. After the first broadcast period (part (b) of the figure,) vehicles 2 and 3 hear vehicle 1's broadcast about itself, and store that information. The same happens for vehicle 4 hearing vehicle 3's broadcast message. After the next broadcast period (part (c) of Figure 1), vehicle 4 hears the message broadcast by vehicle 3 which includes information about all of 1, 2, and 3, and updates its local information. Note that the relative position of the vehicles have changed over time, and some have changed their lanes.

TrafficView does not suffer from memory limitation due to the small size of the stored records. As will be shown in Section 4, the average size for data records is on the order of 50 bytes. Assuming a very high density, five-lane road in which the distance between consecutive vehicles is 5 meters, about 5K bytes will be needed to store the information about all the vehicles in 100 meters, and about 1M bytes to store information of all the vehicles in 20Km. Most of the current portable devices come with more memory than these values.

On the other hand, assuming a transmission range of 250m for the wireless network card, there will be 50 vehicles competing for the same wireless medium in a single lane, and about 250 vehicles in a five-lane road assuming the lanes are close to each other. Hence, the total amount of data that needs to be broadcast by these vehicles every broadcast period is 250MB, which is beyond the capabilities of the current wireless technology. To cope with the bandwidth limitation, each vehicle is allowed to broadcast a small packet—a few kilobytes in size—every broadcast period to allow other surrounding vehicles to share the medium. Therefore, compression/aggregation mechanisms are needed to reduce the size of information to fit into the broadcast packet.

For simplicity, we assume throughout this paper that the road is straight. In the general case, the direction of the movement of a vehicle can be included in the record sent out about that vehicle, and then used to estimate its position on the road trajectory. Moreover, without loss of generality, we assume that the road is along the y axis, and all the vehicles are moving in the positive direction of the road. In a real situation, a road might be bidirectional, where vehicles move in two opposite directions. In this case, a vehicle will need to examine the movement vector in a record received about another vehicle, and ignore it if that vehicle is moving in the opposite direction. This can also be applied in the case of an intersection where a vehicle might hear about different vehicles moving in different directions.

4 System Design

In this section we present the design of the implemented prototype of TrafficView system. Hereafter we use the terms “vehicle” and “node” interchangeably.



Figure 2: Hardware components used in the TrafficView prototype

4.1 Hardware

We implemented a prototype of the TrafficView system as shown in Figure 2. In this prototype, each vehicle is equipped with a portable computer (e.g., Compaq iPAQ with Linux Familiar distribution) augmented with two slots of PCMCIA sleeve, Global Positioning System (GPS), 802.11b wireless network card, DSP-100 2-port RS-232 serial PCMCIA card [1], and an OBDI-II interface [2]. The GPS receiver provides the latitude and longitude of the vehicle in addition to the global time. Using the wireless card, network connectivity is established, and the vehicle is able to send and receive information about other vehicles. The TrafficView software on the node periodically queries the vehicle's status (e.g., speed) using the OBDI-II interface. The DSP-100 card is used to connect the iPAQ to the GPS receiver and the OBD-II interface.

The portable computer runs the TrafficView software that displays road traffic information and is responsible for the interaction with other vehicles. We assume that each vehicle has a unique identification number (ID). The vehicle identification number (VIN) or driver's driving license number are good candidates for this ID.

4.2 Software

In TrafficView, each vehicle stores records about itself and other vehicles it knows about. In this section, we describe the record format and the system modules.

4.2.1 Data Representation

Each record stored about another vehicle consists of the following fields:

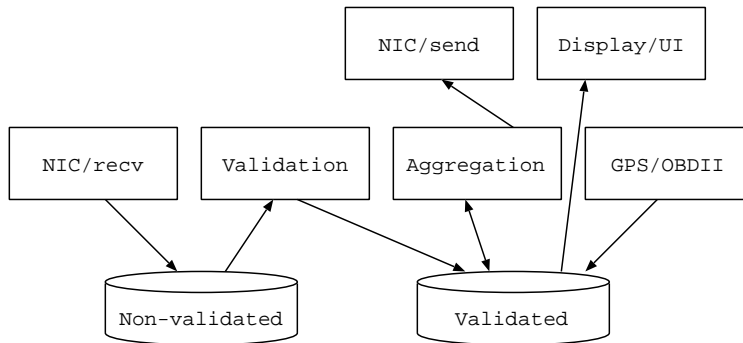


Figure 3: The structure of a node in TrafficView

- *Identification (ID)*: Used to uniquely identify the records belonging to different vehicles.
- *Position (POS)*: The current estimated position of the vehicle.
- *Speed (SPD)*: Used to estimate the vehicle's position if no messages containing information about that vehicle are received.
- *Broadcast Time (BT)*: The global time at which the vehicle broadcast that information about itself.

4.2.2 System Components

Figure 3 shows the software components (modules) of a node in the system. Each vehicle stores records about other vehicles in its local datasets. When the record is first received in a broadcast message, it is stored in the *non-validated* dataset, since it might contain outdated or conflicting information. After these records are examined for validity, they are moved and merged with the *validated* dataset.

A TrafficView node, as shown in Figure 3, contains six modules that operate on its datasets:

- *GPS/OBD module* periodically updates the vehicle's own record in the validated dataset.
- *Receive module* listens to broadcast messages from neighboring vehicles, and stores the records received in the non-validated dataset. It ignores the messages broadcast by its own vehicle.
- *Validation module* validates and resolves conflicts of the records in the non-validated dataset. It then merges the validated versions with the records in the validated dataset. For example, this module removes all the records that are

about vehicles behind its own vehicle¹. Another example of a validity check is when there are multiple records containing information about the same vehicle. In this case, this module keeps the most recent record, and removes the older versions. In addition, this module periodically updates the estimated position of the vehicles in the validated dataset using the stored speeds. The validation module is also responsible of *information aging*, which will be discussed in detail in Section 5.4

- *Aggregation module* performs aggregation algorithms on the records in the validated dataset in order to be able to place more information in the outgoing broadcast messages. This module might as well update the dataset by replacing the original records with the new aggregated version.
- *Send module* writes the contents of the records in the validated dataset in a broadcast message and broadcasts it on the wireless channel using the wireless card.
- *Display/UI module* is responsible of displaying the validated records periodically on the display. It is also responsible for the user interaction (e.g., graphically and/or audibly).

5 Data Aggregation Mechanisms

A MAC layer protocol (e.g., IEEE 802.11b protocol) limits the size of the payload that is sent on the network channel to a maximum size (which is 2312 bytes for 802.11b.) This limit applies to the broadcast message a node intends to send in the TrafficView system as well. On the other hand, the number of records in a node's validated dataset can be large, making it impossible to fit all of them in one broadcast message. In order to deliver as much information about other vehicles as possible, data compression/aggregation techniques should be applied to the validated records. Data compression and aggregation are two different concepts. Data compression is actually "binary compression" in the sense that it does not base the decisions made on the semantics of the data. Moreover, data compression techniques require a lot of computation resources which is not suitable for most portable devices. In this paper we focus on data aggregation mechanisms only.

Data aggregation is based on the semantics of the data. For example, the records from two vehicles can be replaced by a single record with little error, if the vehicles are very close to each other, and they are moving with relatively the same speed; in other words, the distance between them is always in a small range. The way data aggregation contributes to the TrafficView system is by delivering as many records as possible in one broadcast message. This way, more new records can be delivered in a certain period of time and the overall system performance is improved.

¹TrafficView only stores information about the vehicles in front of the current vehicle, and ignores the ones behind it.

5.1 Data Aggregation Basics

A single aggregated record will represent information about a set of vehicles. In this paper we adopt one simple format for the aggregated records²: In an aggregated record, the ID field is extended to a list of vehicles' IDs while the other fields—position, speed, and broadcast time—remain as single values for all the vehicles stored in the record. Formally, if the records $(ID_1, POS_1, SPD_1, BT_1) \dots (ID_n, POS_n, SPD_n, BT_n)$ are being aggregated, and d_i is the estimated distance between the current vehicle and the vehicle with ID_i , the aggregated record will be

$$(\{ID_1, \dots, ID_n\}, POS_a, SPD_a, BT_a)$$

where

$$\begin{aligned} BT_a &= \min\{BT_1, \dots, BT_n\}, \\ POS_a &= \sum_{i=1}^n POS_i/d_i, \text{ and} \\ SPD_a &= \sum_{i=1}^n SPD_i/d_i. \end{aligned}$$

We realize that storing the minimum broadcast time—as opposed to storing the maximum or average—is advantageous, in that it allows the information about the vehicle which corresponds to the minimum broadcast time value to be updated as soon as a fresher record is heard about that vehicle.

According to the way the aggregated fields are calculated, the aggregated records should have close values for their POS, SPD, and BT fields to reduce the error resulting from the aggregation. Our preliminary experiment results showed (as expected) that if two vehicles are close to each other, their broadcast records will arrive at a specific receiver at nearly the same time. Figure 4 shows the average difference between the record broadcast time and its receipt time, and the distance between the sender and the receiver, for a simulation of 550 total nodes, moving with an average speed of 30m/s, using the simple diffusion mechanism for information exchange. As a result, if two records have close POS values, they are expected to have close BT values.

At the same time, if the difference between the speed of two vehicles that are close to each other is big, their distance will grow in a short time as well. Keeping in mind that the broadcast period is in the order of seconds, we can ignore the speed difference among the aggregated records, because the record will be updated with the new up-to-date position information as soon as new broadcast messages are heard. As a conclusion, the records are selected for aggregation based of their relative distances only. To achieve this in an efficient manner, records are kept sorted on the estimated relative distance of the current vehicle to the corresponding vehicles. This way, we limit the search for records to combine only to consecutive records in the sorted list.

Whenever a node receives a record containing information about some vehicles, it first checks the information in that record against the validated records it has. If the

²We are developing other aggregation formats for the TrafficView system.

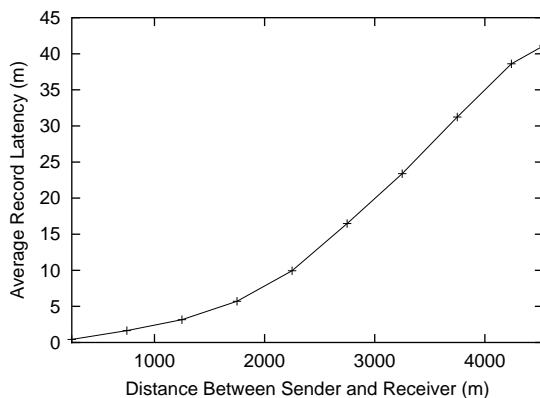


Figure 4: Average record delay based on the distance between the sender and receiver

record contains information about some vehicles about which the node already knows, it performs the following:

1. If the broadcast time of the records is greater than the broadcast time of the stored record, it means the new record is fresher, and therefore the node removes the corresponding vehicle ID from its stored record,
2. Otherwise, the new record contains older information, and hence the node removes the corresponding vehicle ID from the received record.

In TrafficView, vehicles apply the aggregation procedure on the records in the validated dataset each broadcast period to prepare the broadcast packet. Our preliminary experiments showed that the effect of each vehicle either replacing its current validated records with the aggregated version, or maintaining the original records in its validated dataset, on the quality of the information gained by other vehicles on the road, is almost identical; the only difference being the imposed overhead in the next broadcast period. We therefore decided to replace the validated dataset records with the new aggregated version during each broadcast period in order to reduce the overall aggregation overhead.

In the following subsections, we describe different algorithms to select the aggregated records. Table 1 lists a set of records that will be used for the illustration of the algorithms.

5.2 Ratio-based Algorithm

The algorithm divides the road in front of the vehicle to a number of regions (r_i). For each region, an aggregation ratio (a_i) is assigned. The aggregation ratio is defined as the inverse of the number of individual records that would be aggregated in a single record. Each region is assigned a portion (p_i where $0 < p_i \leq 1$) of the remaining free space in the broadcast message. The aggregation ratios and region portion values are

Algorithm 1: RATIO-BASED ALGORITHM()

INPUT :

Sorted list of validated records
 n : number of regions ($r_1 \dots r_n$)
 $a_1 \dots a_n$: aggregation ratios
 $p_1 \dots p_n$: message portion values

OUTPUT :

$th_1 \dots th_n$: merging thresholds
 $b_1 \dots b_n$: region boundaries

VARIABLES :

R : size of the remaining space in the broadcast message
 L : number of records left in the list of records
 $optimum$: optimum aggregation ratio
 d_{max} : distance of the farthest vehicle the current vehicle knows about
 l_i : number of records in region i

ALGORITHM :**main**

Initialize b_i and th_i to 0 for all i

$b_0 \leftarrow d_{max}$

$R \leftarrow$ size of broadcast message

$L \leftarrow$ number of records in the input list

for each region r_i

$\left\{ \begin{array}{l} optimum \leftarrow \frac{R}{(\text{average record size}) \times L} \\ \mathbf{if} \ optimum \geq 1 \\ \quad \mathbf{then} \ return \\ \mathbf{if} \ optimum \geq a_i \\ \quad \mathbf{then} \ \left\{ \begin{array}{l} b_i \leftarrow d_{max} \\ th_i \leftarrow \frac{b_i - b_{i-1}}{L \times optimum} \\ \end{array} \right. \\ \quad \mathbf{return} \\ \mathbf{do} \ \left\{ \begin{array}{l} l_i \leftarrow \text{number of records that fit in} \\ \quad \frac{R \times p_i}{a_i} \text{ bytes} \\ L \leftarrow L - l_i \\ \mathbf{if} \ l_i = 0 \\ \quad \mathbf{then} \ \left\{ \begin{array}{l} b_{i-1} \leftarrow d_{max} \\ \end{array} \right. \\ \quad \mathbf{return} \\ b_i \leftarrow \text{relative distance of the last record fit} \\ th_i \leftarrow \frac{b_i - b_{i-1}}{l_i \times a_i} \\ R \leftarrow R - R \times p_i \end{array} \right. \end{array} \right.$

ID	relative distance	speed	broadcast time
1	40	30	9.80
2	65	25	9.75
3	120	35	9.00
4	140	20	8.80
5	250	30	6.90
6	280	15	6.75
7	600	30	4.25

Table 1: Sample records used to illustrate different aggregation algorithms

assigned according to the importance of the regions and how accurate the broadcast information about the vehicles in that region is needed to be. For example, assigning decreasing values to the aggregation ratios and equal values to portion parameters will result in broadcasting less accurate information about regions that are farther away from the current vehicle, since for those regions, each individual record will represent large number of aggregated vehicles (records).

Given the aggregation ratios, portion values, and number of regions, the algorithm calculate the region boundaries ($[b_i, b_{i+1}[$) as shown in Algorithm 1. Knowing the number of current records in the validated dataset that lie within the boundaries of each region and the corresponding free space in the broadcast packet, the algorithm calculates the merging threshold (th_i) corresponding to each region. Any set of consecutive records in region r_i will be aggregated in a single record if the relative distance (in y direction) between the first and the last record is less than the corresponding merge threshold, th_i .

As shown in Algorithm 1, the algorithm will not over-aggregate the records. This is guaranteed by calculating the optimum aggregation ratio at the beginning of the loop for each region. This aggregation ratio is the value needed to fit the rest of the records in the message free space. If this ratio is greater than or equal to one, the algorithm terminates since no aggregation is needed. Otherwise, the optimum value and the aggregation ratio of the current region are compared and the maximum among these two is used.

After the algorithm aggregates the records, it starts writing the record contents to the broadcast message until no free space is left. There is no guarantee to write all the record contents in the message. The tradeoff between the number of records written and the accuracy of the records is governed by the parameter values used in the algorithm.

As an example, assume a vehicle with $ID = 0$, using this algorithm, divides the road into two regions, and the corresponding parameter values are $a_1 = 0.5$ with $p_1 = 0.5$ and $a_2 = 0.25$ with $p_2 = 0.5$. If the algorithm is applied to the records listed in Table 1, it will calculate the following parameters: $b_1 = 120$, $th_1 = 80$, $b_2 = 600$, and $th_2 = 261.82$. Note that th_2 is calculated using the optimal aggregation ratio 0.46 instead of the input value, 0.25.

After calculating the parameters, in the first region, the algorithm first combines records 1 and 2, and then combines the result with record 3. Likewise, the records 4,

ID(s)	relative distance	speed	broadcast time
1, 2, 3	61.58	29.30	9.00
4, 5, 6	203.88	21.73	6.75
7	600	30	4.25

Table 2: Records sent out by the Ration-based algorithm

5, and 6 are combined in the second region. The records sent out by the algorithm are shown in Table 2. Record 7 is sent not aggregated. Note that the speeds are calculated as the weighted average in a single value while the broadcast times are calculated as the minimum value.

5.3 Cost-based Aggregation

In the Ratio-based algorithm, records that satisfy the merging threshold, (th_i), criterion are “blindly” combined without considering the cost of the aggregation. In contrast, the Cost-based algorithm assigns a cost for aggregating each pair of records, and whenever it needs to aggregate two records, the two that correspond to the minimum cost are chosen. Assume two records storing aggregated information about s_1 and s_2 vehicles, with a relative distance of d_1 and d_2 , respectively. The cost of aggregating the two records is calculated as follow:

$$cost = \frac{|d_1 - d_a| \times s_1 + |d_2 - d_a| \times s_2}{d_a},$$

where d_a is the relative distance of the aggregated group of records (vehicles). This formula is calculated such that it

1. assigns a high cost for the vehicles that are relatively close to the current vehicle ($1/d_a$),
2. tries to minimize the error introduced during the merging ($|d_i - d_a|$), and
3. minimizes the number of vehicles affected by the aggregation (s_i).

The details of the algorithm are shown in Algorithm 2. The aggregation ratios and message portion values are the inputs to the algorithm. For each aggregation ratio and the corresponding portion value, the algorithm starts by continuously selecting the two records that result in the minimum cost, and aggregating them until the number of records is reduced to the value needed by the factor of the aggregation ratio. Afterwards, it writes the contents of the first records in the sorted list to the beginning of the message until they fill the space allocated according to the corresponding portion value. In the next iteration, the same procedure of aggregation and writing is applied to the rest of the records that are not written yet. The aggregation ratios in each iteration is compared with the optimum aggregation ratio to avoid over-aggregation.

A problem that might happen is that as the algorithm proceeds, the number of records left decreases, and the distance between any two consecutive records increases.

Algorithm 2: COST-BASED AGGREGATION()

INPUT :

Sorted list of validated records
cost-threshold
 n : number of regions ($r_1 \dots r_n$)
 $a_1 \dots a_n$: aggregation ratios
 $p_1 \dots p_n$: message portion values

VARIABLES :

R : size of the remaining space in the broadcast message
 L : number of records left in the list of records
optimum : optimum aggregation ratio
 l_i : number of records in region i

ALGORITHM :**main**

$R \leftarrow$ size of broadcast message
 $L \leftarrow$ number of records in the input list
for each region r_i
 $\left\{ \begin{array}{l} \textit{optimum} \leftarrow \frac{R}{(\textit{average record size}) \times L} \\ \textit{if } \textit{optimum} \geq 1 \\ \quad \textit{then return} \\ \textit{a}_i \leftarrow \max(\textit{optimum}, \textit{a}_i) \\ \textit{goal} \leftarrow \textit{a}_i \times L \\ \textit{while } L > \textit{goal} \\ \quad \textit{do } \left\{ \begin{array}{l} \textit{c} \leftarrow \textit{minimum cost of merging two} \\ \quad \textit{consecutive records in the remaining} \\ \quad \textit{records set} \\ \quad \textit{if } \textit{c} > \textit{cost-threshold} \\ \quad \quad \textit{then return} \\ \quad \textit{Merge the two records corresponding to the} \\ \quad \quad \textit{minimum cost} \\ \quad \quad L \leftarrow L - 1 \end{array} \right. \\ \textit{l}_i \leftarrow \textit{number of records that fit in } R \times \textit{p}_i \textit{ bytes} \\ R \leftarrow R - \textit{size of the } \textit{l}_i \textit{ records} \end{array} \right.$

ID(s)	relative distance	speed	broadcast time
1, 2	49.52	27.84	9.75
3, 4	129.23	27.80	8.80
5, 6	264.15	22.72	6.75

Table 3: Records sent out by the Cost-based algorithm

Hence there is a risk of combining two records that correspond to vehicles that are too far away from each other. To avoid this problem, the algorithm terminates as soon as the calculated cost is greater than a threshold parameter (*cost-threshold*.)

For example, assume vehicle with $ID = 0$ intends to use this algorithm for the records listed in Table 1, where $a_1 = a_2 = 0.5$, $p_1 = p_2 = 0.5$, and *cost-threshold* = 0.9. During the first iteration (a_1), it first aggregates records 5 and 6 (cost = 0.11), then 3 and 4 (cost = 0.15), and finally 1 and 2 (cost = 0.50). In the second phase (a_2), the minimum cost is 1.22, which is greater than the cost threshold, therefore the algorithm terminates. Table 3 lists the records that are sent out by vehicle 0 and the corresponding fields. In this case, vehicle 0 cannot fit record 7 in its message.

5.4 Information Aging

The records stored in both the validated and non-validated datasets, must be examined to verify that they reflect the current state of the road and eliminate the outdated (old) information. For example, vehicles included in the validated dataset might have exited the road. Moreover, new received records (non-validated) might contain inaccurate information due to frequent changes in the speed of the corresponding vehicles and/or aggregation mechanisms applied to the data within relaying nodes.

There are two main problems here: how should the value of the information in a broadcast message be assessed, and how can a balance between knowing inaccurate information about a vehicle, and having no knowledge about it, be achieved. In general, if the cost of knowing inaccurate information about vehicle j that is at a relative distance of d is a function $c_1(j, d)$, and the cost of having no information about j is another function $c_2(j, d)$, the information should be accepted and stored if $c_1(j, d) < c_2(j, d)$, otherwise it should be dropped. Unfortunately, it is not clear how to assign values to these two functions.

To solve this problem, TrafficView exploits two aging mechanisms. The first mechanism associates a timer with each record added to the validated dataset. This timer is reset each time the record is updated due to a broadcast message. If the timer is expired, the record is dropped. The second mechanism, which we call Receive-aging, deals with newly received records via broadcast messages. Whenever a new record is received, the expected latency in receiving the record is calculated and compared to the actual latency (the difference between the receive time and the BT field.) If the difference between these two is lower than a threshold, it is stored; otherwise, it is considered out-of-date, and is ignored.

Formally, assume node 2 receives a record about vehicle 1 at time t . Looking at the

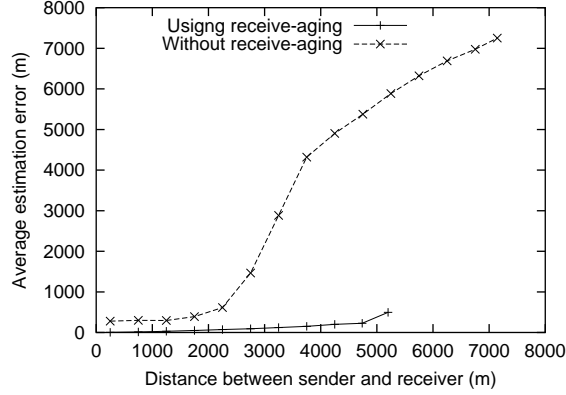


Figure 5: Effect of Receive-aging: average position estimation error in to simulations, one with Receive-aging enabled, and the other without that mechanism

record contents, node 2 extracts the time BT_1 at which the record was first broadcast, and vehicle 1's position POS_1 at that time. Knowing its own position POS , node 2 estimates its position POS_2 at time BT_1 as

$$POS_2 = POS - v_2 \times (t - BT_1),$$

where v_2 denotes node 2's speed which we assume, with no loss of generality, to be fixed during the time period $[BT_1, t]$. Node 2 then calculates the expected delay in receiving the record as:

$$delay = \frac{|POS_1 - POS_2|}{|r/p + v_2|}.$$

where r is the wireless transmission range, and p is the broadcast period. Therefore, r/p is the approximate propagation speed of the information between the vehicles. This record is then accepted by node 2 only if

$$|t - BT_1| \leq \delta_1 + (1 + \delta_2) \times delay,$$

where δ_1 and δ_2 are acceptance thresholds.

To validate the effectiveness of the Receive-aging mechanism, we ran two simulations with 870 total nodes moving with an average speed of 30m/s. In the first run, the nodes were using this mechanism with $\delta_1 = 6.0$ and $\delta_2 = 0.3$, whereas in the second, it was disabled. Figure 5 presents average estimation error of the position of the vehicles in the two runs for different distance between the sender and receiver. As the figure shows, when Receive-aging is not used, the estimation error for vehicles at far away distances is huge. In contrast, using this mechanism has reduced the average error to a small value.

6 Performance Evaluation

We have implemented the algorithms and have run ns-2 simulations to compare the performance of different algorithms. In this section, we present the experiments, and the corresponding results.

6.1 Scenario Generator

Modelling road traffic is a research topic about which a lot of work has been done. For example CORSIM [6] is a microscopic traffic simulator developed by The Federal Highway Administration. Unfortunately, none of the traffic modeler tools are freely available to public. We have therefore developed our own scenario generator tool based on “setdest”—a generator tool for random-way point mobility model, developed at Carnegie Mellon.

The scenario generator accepts as parameters simulation time, length of the road, average speed of the nodes, number of lanes on the road, and the average gap length between vehicles. It uses a simplified traffic model as follows:

- *Entries and Exits:* The entries and exits are evenly distributed along the road each 1000 meters. Vehicles may enter the road at each entry except the last one and leave at any consecutive exit. Vehicles enter the road at the front-end entry with a probability of 0.7, and at side entries with a probability of 0.3.
- *Speed Changes:* To model the changes to the speed of a node, the road between the entry point and exit point of a node is divided into regions of 50 meters, and a constant speed of $\text{max speed} \times (0.75 + \text{rand}(-2, 2) \times 0.125)$ is used for each region, where $\text{rand}(a, b)$ returns a uniformly distributed random integer between a and b .
- *Changing Lanes:* Vehicles can change their lanes with no dependence on other vehicles. The probability of staying on the same lane is 0.6 whereas the probability of changing to the right or left lane is 0.2. If a node is on the leftmost (rightmost) lane, the probability of changing to the left (right) lane is 0.0.
- *Vehicle Density:* The density of vehicles is an important factor because it determines the number of neighboring nodes in the transmission range of a vehicle, which has a great impact on the transmission delay and available bandwidth of the network. The scenario generator initially puts

$$\frac{\text{road-length} \times \text{number of lanes}}{\text{average gap}}$$

active nodes, evenly distributed, on the road. Once a vehicle leaves the road at one of the exits, it is deactivated, and a new node is added (activated) to the road randomly. As soon as a node is deactivated, it will no longer affect our metric calculations introduced in the next section.

Figure 6 shows the histogram of the average speed and number of lane changes per minute for a scenario generate with average speed = 30m, and average gap = 100.




Figure 6: Sample histograms of average speed (left) and average number of lane changes per minute (right) in a scenario generated by the scenario generator tool

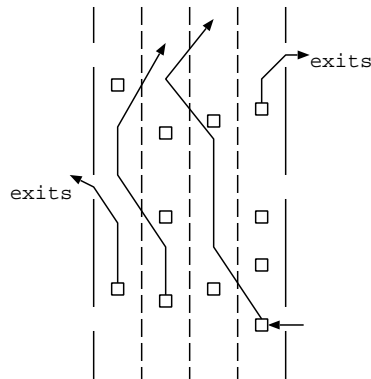


Figure 7: A segment of a road in an example scenario generated by the scenario generator

The graphs show the percentage of vehicles that have that average speed and average number of lane changes per minute, respectively. A segment of a road in an example scenario generated by the tool is shown in Figure 7. The road, along which 11 nodes are moving, has three exits at each side.

For all the simulations in this paper, we fixed the length of the road to be 15,000 meters with 4 lanes. We used 802.11b (with a data transmission rate of 11Mb) as the wireless media with a transmission range of 250m. During a simulation, nodes broadcast messages periodically. The broadcast period is selected uniformly from [1.75, 2.25] seconds, and each node recalculates the next broadcast period after the current broadcast. For all the simulation runs, we use broadcast messages of size 2312 (the maximum payload size of 802.11b standards) and we fix the simulation time to 300 seconds.

6.2 Algorithms and Metrics

We implemented two simple algorithms in addition to the ones introduced in Section 5 for comparison purposes: non-aggregation and brute-force cost-based. In the non-aggregation method, no aggregation is performed and each node broadcasts only the first records in its validated dataset that fit in one broadcast message. In the brute-force cost-based algorithm, the node keeps aggregating its records using the same technique introduced in the Cost-based algorithm, until it can fit all the current record contents in one broadcast message.

During the simulation, each node periodically broadcasts the information it has about other vehicles. Upon receipt of a broadcast message, a node updates its datasets accordingly.

We will use the following metrics and graphs to assess the performance of the algorithms:

- *Accuracy*: The road in front of each vehicle is divided into regions of 500 meters long, and the average error in estimating the position of vehicles in each region is calculated. In the accuracy graphs, the average estimation error for each region is shown, averaged over all the nodes during the corresponding simulation.
- *Visibility*: We define the visibility of a specific vehicle as the *average* relative distance to the vehicles it knows about. A point (d, p) on a visibility graph means that $p\%$ of the vehicles have had a visibility of d meters or more.
- *Knowledge Percentage*: The road in front of each vehicle is divided into regions of 200 meters long. For each region, the percentage of the vehicles in that region about which the current node knows, is defined as the knowledge percentage of that node for that region. The knowledge percentage graph presents the knowledge percentage for each region, averaged over all the nodes during a simulation run.

6.3 Aggregation Parameters

We ran different simulations to select the suitable values for the parameters of the Ratio-based and Cost-based algorithms. Total number of nodes in all of these simulations was 960. The suitable set of values are used in the simulations run to compare the performance of different algorithms.

For these two algorithms, the maximum number of regions in front of each node is four. The first three regions are defined by parameters a_1, a_2, a_3, p_1, p_2 and p_3 . The fourth region is defined dynamically by the remaining available space in the outgoing message and the remaining set of records that each node has.

Table 4 lists the parameters used in different runs of the algorithms. The way these parameters are selected is to first run the algorithm with param1, and param2 to select the better a_i values and then fix a_i s and run with param3 and param4 to choose p_i values. The incentive is to select a_i s as small as possible to achieve as large visibility as possible while maintaining a good accuracy for the closer vehicles.

In param3, we allocate $0.4, (1 - 0.4) \times 0.6 = 0.36$ and 0.192 of the broadcast message to the first three regions, respectively, hence leaving only $0.048 (\times 2312 =$

Name	a_1	a_2	a_3	p_1	p_2	p_3
param1	0.5	0.25	0.17	0.5	0.5	0.5
param2	0.75	0.5	0.25	0.5	0.5	0.5
param3	0.5	0.25	0.17	0.4	0.6	0.8
param4	0.5	0.25	0.17	0.3	0.43	0.75

Table 4: Parameter settings for different runs of the Ratio-based and Cost-based aggregation algorithms




Figure 8: Visibility graphs for different runs of the Ratio-based algorithm

111 bytes) of the message to the rest of the records. In contrast, param4 allocates approximately 0.3 of the message to each of the first three regions and leaves 0.1 of the space to the remaining records. The reason we started with the a_i values is that they have a larger effect on the performance of the aggregation algorithms than the effect of p_i parameters.

Figure 8 shows the visibility graph for different runs of the Ratio-based algorithm. We found out that param1 settings give a higher accuracy while maintaining a good visibility. We therefore use param1 values to set the Ratio-based parameters in the rest of the simulation runs.

On the other hand, we noticed that using param4 gives a higher accuracy among the other settings for the Cost-based aggregation algorithm while maintaining a good visibility as shown in Figure 9. We therefore use the values of param4 in the rest of the simulation runs of the Cost-based algorithm.

For the Receive-aging mechanism, we set δ_1 to 6.0 and δ_2 to 0.3. These values were selected by running the non-aggregation method with different values for these parameters, and choosing the ones that resulted in the best visibility while maintaining an acceptable accuracy.

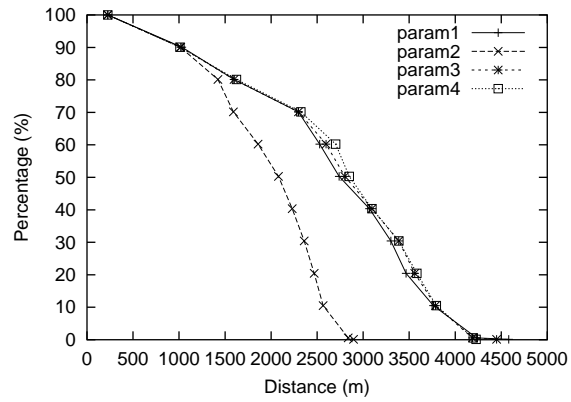


Figure 9: Visibility graphs for different runs of the Cost-based algorithm




Figure 10: Visibility graphs for different runs of the non-aggregation algorithm

Name	Total # of nodes	Average speed	Average gap
scen1	690	10	100
scen2	780	20	100
scen3	870	30	100
scen4	548	40	175

Table 5: Parameters of different simulations used to compare different algorithms




Figure 11: Visibility graphs for different runs of the brute-force algorithm

6.4 Results

To compare the performance of different algorithms, we ran each algorithm on different scenarios. Table 5 lists the parameters of each simulation.

We first look at the effect of the road parameters on different algorithms. Figure 10 shows the visibility graph for runs on different scenarios of the non-aggregation algorithm. We notice in this Figure that average speed does not have a significant effect on the performance of the algorithm. On the other hand, the average gap, directly effects the performance: As the gap between vehicles increases, the number of vehicles scattered over the road decreases. Therefore, the broadcast message will contain records about vehicles in farther distances and that results in increasing the visibility metric.

Figure 11 shows the same graph for the brute-force algorithm. For this algorithm, as the average speed increases, the rate the vehicles get closer to or depart from each other increases. Due to this increase, more number of records get aggregated. With the increase in cars speed, the values of broadcast fields (*BT*) fields decrease faster and that result in invalidating records more quickly due to aging mechanisms, and hence the average visibility decreases. Again, increasing the gap value increases the vehicles visibility. The other aggregation mechanisms show a similar behavior.

We will use the results of running the aggregation algorithms using scenario scen3

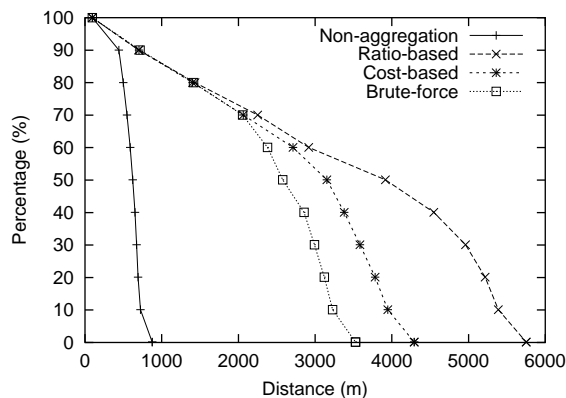


Figure 12: Visibility graphs for runs of different algorithms using scenario scen3

to compare their performance, because it reveals the behavior of the algorithms more clearly.

Figure 12 shows the visibility graph of the different algorithms. The Ratio-based algorithm achieved the highest visibility value. The Cost-based algorithm outperforms the brute-force algorithm. As mentioned earlier, this is due to the fact that records are invalidated more quickly in the brute-force algorithm. The reason the Ratio-based achieves the highest visibility is that it performs aggregation on all the vehicles in all the regions while the Cost-based and brute-force methods have less or no control on selecting the region where the aggregation is performed. The result is that the boundaries of the regions generated by Ratio-based algorithm cover larger road areas than the other algorithms, and hence it has the highest visibility.

Figures 13 and 14 present average estimation error and average knowledge percentage for different algorithms using scen3. As a result of the Ratio-based mechanism performing aggregation on all the regions, its knowledge percentage about the close and medium-distanced vehicles is less than the other algorithms; its accuracy is also lower than the other algorithms.

From the above results we conclude that the Ratio-based algorithm is more flexible than the other algorithms in that it provides more control over the tradeoff between the accuracy and visibility governed by the parameter setting. For the other methods, although tuning the parameters is easier, the cost function does not provide the flexibility present in the Ratio-based algorithm.

7 Related Work

A lot of effort has been put into Inter-Vehicle Communication research in the recent few years. CarNet [12] project focuses on how the radio nodes in the vehicles get IP connectivity with the help of Grid [9]. A Grid-to-Internet gateway is setup to route packets between the wireless network and the Internet. In [14], a wireless traffic light




Figure 13: Average position estimation error for runs of different algorithms using the scen3 scenario




Figure 14: Average knowledge for runs of different algorithms using the scen3 scenario

system is presented. At the intersection, a static control unit periodically broadcasts the current light status, location of intersection, and a reference point, using which the vehicles approaching the intersection can check their relative position and make a decision accordingly. They also designed collision warning system [11] in which peer-to-peer beacon message exchange is used. An architecture of the vehicular communication is described in [5]. It integrates inter-vehicle communication (IVC) with Vehicle-Roadside Communication (VRC), where both moving vehicles and base stations can be peers in the system. The peers are organized into Peer Spaces for message exchange, in which flooding is the main method of delivery.

The authors of [13] examine the feasibility of short range communication between fast moving vehicles using Bluetooth, and a mobile test-bed RUSH has been established in [17], composed of the fixed base station and mobile nodes on shuttle buses.

Two delivery modes known as pessimistic and optimistic forwarding are compared in disconnected vehicle networks in [4]. The experiment shows that the average delay in optimistic delivery is better. The authors of [3] propose a "wait-and-resend" scheme where a mobile node can cache the message for a while before new neighbors enter its transmission range, and [10] proposes an algorithm to dynamically modify the trajectories of the intermediate nodes to approach next available nodes, for relaying the message to the destination.

8 Conclusions and Future Work

In this paper we introduced the TrafficView system, which is a part of broader project—e-Road—that is still under development. The goal of TrafficView is to provide the driver of a vehicle with information about traffic and road conditions. The essence of the system is to gather and disseminate traffic information between the vehicles on the road. We presented the basic design of the system, and the algorithms used for data aggregation and information dissemination using the 802.11b wireless standards.

For future work, we are continuing to work in a number of different directions. We are experimenting with a linear programming model to estimate the aggregation parameters dynamically based on the road condition.

Privacy is an important issue in such a system. Different privacy levels should be available from which the drivers can select. One level of privacy could be to completely hide any information about the vehicle while it continues to participate in relaying other vehicles' information. Another level is to allow others to gain information about the vehicle without being able to identify it. For example, vehicles on the road may know about a group of vehicles, including the target one which is located in an area on the road without being able to identify the exact location of any vehicle including the target.

Security and trust are two other important issues in such a system. A fraudulent vehicle could disseminate information about nonexistent vehicles, or broadcast bogus information about existing vehicles. Different mechanisms should be proposed to prevent this and to identify those fraudulent vehicles and avoid accepting their packets.

We believe that TrafficView and the e-Road project will greatly enhance and ease the driving experience. At the same time, they will encourage and trigger several applications to be built over these systems.

References

- [1] <http://www.quatech.com>
- [2] <http://www.obddiagnostics.com>
- [3] L. Briesemeister, G. Hommel, "Role-based multicast in highly mobile but sparsely connected ad hoc networks," in *First Annual Workshop on Mobile Ad Hoc Networking and Computing*, pp. 45–50, Aug. 2000.
- [4] Z. D. Chen, HT Kung, D. Vlah, "Ad Hoc Relay Wireless Networks over Moving Vehicles on Highways," in *Proceeding of the 2001 ACM ACM International Symposium on Mobile Ad-hoc Networking and Computing*, pp. 247–250, Oct. 2001.
- [5] I. Chisalita, N. Shahmehri, "A peer-to-peer approach to vehicular communication for the support of traffic safety applications," *5th IEEE Conference on Intelligent Transportation Systems*, pp. 336–341, Singapore, Sept. 2002.
- [6] CORSIM User Manual, Version 1.01, The Federal Highway Administration, US Department of Transportation, 1996.
- [7] General Motors Collaborative Laboratory website available online at <http://gm.web.cmu.edu/>.
- [8] W. Kellerer, "(Auto)Mobile Communication in a Heterogeneous and Converged World," *IEEE Personal Communications*, Vol. 8(6), pp. 41–47, Dec. 2001.
- [9] J. Li, J. Jannotti, D. S. J. De Couto, D. R. Karger, R. Morris, "A Scalable Location Service for Geographic Ad Hoc Routing", *ACM Mobicom 2000*, pp. 120-130, Boston, MA.
- [10] Q. Li, D. Rus, "Sending Messages to Mobile Users in Disconnected Ad-hoc Wireless Networks," in *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking*, Boston, MA, Aug. 2000.
- [11] R. Miller, Q. Huang, "An Adaptive Peer-to-Peer Collision Warning System," *IEEE Vehicular Technology Conference (VTC)*, Birmingham, AL, May 2002.
- [12] R. Morris, J. Jannotti, F. Kaashoek, J. Li, D. Decouto, "CarNet: A Scalable Ad Hoc Wireless Network System," *9th ACM SIGOPS European Workshop*, Kolding, Denmark, Sept. 2000
- [13] P. Murphy, E. Welsh, P. Frantz, "Using Bluetooth for Short-Term Ad-Hoc Connections Between Moving Vehicles: A Feasibility Study," *IEEE Vehicular Technology Conference (VTC)*, pp. 414–418, Birmingham, AL, May 2002.
- [14] Q. Huang, R. Miller, "The Design of Reliable Protocols for Wireless Traffic Signal Systems". Technical Report WUCS-02-45, Washington University, Department of Computer Science and Engineering, St. Louis, MO.

- [15] M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *IEEE Personal Communications*, Vol. 8(4), pp. 10–17, Aug. 2001.
- [16] C. Schwingenschloegl, T. Kosch, "Geocast Enhancements for AODV in Vehicular Networks," *ACM Mobile Computing and Communications Review*, Vol. 6(3), pp. 96–97, Jul. 2002.
- [17] E. Welsh, P. Murphy, P. Frantz, "A Mobile Testbed for GPS-Based ITS/IVC and Ad Hoc Routing Experimentation," *International Symposium on Wireless Personal Multimedia Communications (WPMC)*, Honolulu, HI, Oct. 2002.