

On the Effectiveness of an Opportunistic Traffic Management System for Vehicular Networks

Ilias Leontiadis*, Gustavo Marfia^{††}, David Mack*,
Giovanni Pau^{†§}, Cecilia Mascolo* and Mario Gerla[†]

*Computer Laboratory, University of Cambridge, UK

†Computer Science Department, UCLA, CA, USA

‡Computer Science Department, University of Bologna, Italy

§Ubiquitous Computing Laboratory - Macau Polytechnic Institute, Macau SAR.

Abstract—Road congestion results in a huge waste of time and productivity for millions of people. A possible way to deal with this problem is to have transportation authorities distribute traffic information to drivers, which in turn can decide (or be aided by a navigator) to route around congested areas. Such traffic information can be gathered by relying on static sensors placed at specific road locations (e.g., induction loops, video cameras) or by having single vehicles report their location, speed and travel time. While the former approach has been widely exploited, the latter has seen birth only more recently, and, consequently, its potential is less understood. For this reason, in this paper we study a realistic test case that allows to evaluate the effectiveness of such a solution. As part of this process: a) we designed a system that allows vehicles to crowd-source traffic information in an Ad-Hoc manner, allowing them to dynamically reroute based on individually collected traffic information, b) we implemented a realistic network-mobility simulator that allowed us to evaluate such a model, and c) the main focus of this paper: we performed a case study that evaluates whether such a decentralized system can help drivers to minimize trip times. This study is based on traffic survey data from Portland, Oregon and our results indicate that such navigation systems can indeed greatly improve traffic flow. Finally, to test the feasibility of our approach we implemented our system and run some real experiments at UCLA’s C-Vet test-bed.

Index Terms—Ad-Hoc, Wireless Communication, Intelligent Transportation System (ITS), Realistic Traffic Flow Simulation, Vanet

I. INTRODUCTION

According to “The 2007 Urban Mobility Report” published by the Texas Transportation Institute, in year 2005 the traffic congestion cost to the nation (in the 437 US urban areas) added up to \$78.2 billion. The same report also observes that more and more time is spent in cars due to congestion. A measure of this phenomenon is the increase of the Travel Time Index (TTI), the ratio of travel time in rush hours to travel time at quiet periods that has steadily risen from 1.09 in 1982 to 1.26 in 2005. This trend has led to the development of Intelligent Transportation System (ITS) technologies. Although there is a general consensus about the role ITSs may play in reducing traffic, such systems are not necessarily guaranteed to succeed in reducing congestion.

During the last few years, a wealth of research has studied how to optimize the internetworking of vehicles utilizing

short-range radios (e.g., WiFi, DSRC), in order to support a wide range of services. In particular, Intelligent Transportation Systems could benefit from the use of Vehicular Ad-Hoc Networks (VANET) in metropolitan areas by enabling each vehicle to act as a traffic probe that measures and afterwards spreads traffic related information. This information can then be used by other vehicles to efficiently select their routes in order to avoid congested areas. In essence, an Intelligent Transport System can use the vehicles themselves to crowd-source traffic information.

Clearly, the advantages deriving from the deployment of a VANET-based ITS are manifold. Many more streets, than those nowadays equipped with a monitoring infrastructure, could be easily observed without requiring additional costs (currently only major urban areas can afford monitoring infrastructure). Furthermore, many more services that may utilize the same infrastructure (e.g., pollution management, accident prevention, etc.) could be provided, thus requiring only limited additional investments.

While cellular networks can be used to offer some of these services [1], this solution can also create a number of issues. First of all, the service providers in each country impose different rules and restrictions as to what kind of data can be exchanged through their network or even what type of applications can access it, making it impossible for vehicular applications to be deployed globally (e.g., at least a per country agreement will be required). Additionally, the cost of cellular data communication is restrictively high, as it can reach a few pence per KB. Even expensive “unlimited” plans are usually capped to a few hundred megabytes per month, making large-scale communication (such as real-time, fine-grained traffic information) between millions of vehicles unfeasible. Furthermore, although 3G connections can support up to 128 Kbits/sec inside a moving vehicle, the bandwidth is shared between all users inside the cell. Even today, when 3G is not widely used, the network is swamped by traffic, resulting in very low throughput in densely populated areas. Finally, Ad-Hoc connectivity is more sensible in order to disseminate local information. The use of one-hop local Ad-Hoc communication does not require any kind of license or any infrastructure deployment and, most importantly, local wireless communication between vehicles is becoming a reality: there is increasing

industrial interest and support for vehicular networks [2], [3], led mainly by the interest to maximize road safety (i.e., to avoid vehicle collisions).

However, before being implemented, VANET-based ITS systems require further assessments, as their beneficial effects on traffic conditions are not yet fully understood. In fact, in a fully decentralized ITS, each vehicle bases its own traffic knowledge and routing decisions on only partial data, as only part of the generated traffic information could be received. As a result, many new research initiatives have taken place with the scope of understanding the impact of distributed ITSs on urban congestion [4].

In this context, we would like to examine the effects of such a distributed system on traffic. To conduct this study we designed a distributed ITS and an evaluation platform based on two realistic simulators, i) a vehicular mobility and ii) a network simulator, that interact and depend, at each step, one on the output of the other. Moreover, the key contribution of this work is that we used these tools in order to examine the impact of this system in a realistic city scenario. Finally, we implemented and ran a small-scale proof-of-concept field tests at UCLA's C-Vet test-bed

Summarizing, with this paper we aim at understanding the impact of a distributed ITS system in a realistic scenario. To the best of our knowledge, this work is the first to consider a real city topology and realistic traffic flows in evaluating such systems. *In fact, origin-destination pairs of all vehicles utilized in our simulations are derived from large-scale surveys on the population of Portland, Oregon* [5]. These data allow us to study a realistic case-study where traffic patterns are matching real observations. Our results show that such systems can be beneficial for minimizing traffic conditions in metropolitan areas, but also open new interesting research directions, as we show that its improvements on traffic vary as a function of different parameters (e.g., flow intensities, traffic information aggregation algorithms, percentage of malicious vehicles, etc).

The remainder of this paper is organized as follows. We start by describing the existing related work in Section II and our main motivation scenarios in Section III. In Section IV we illustrate CATE, our distributed ITS system. In Section V we present our evaluation platform: a simulation platform that combines a mobility and a network simulator. In Section VI we present the extensive study on how such an Ad-Hoc system could affect traffic conditions in a realistic scenario, while Section VII discusses a few simple experimental trials we have conducted to test the feasibility of this solution. Finally, Section VIII concludes the paper listing possible future directions.

II. RELATED WORK

An extensive body of literature exists on traffic congestion causes and effects. Therefore, we only directly cite those works that are particularly relevant in the present study.

First of all, there was a lot of research about how perfect traffic information does not guarantee lower congestion. In fact, assuming all vehicles were provided with real-time traffic information about every road segment and each vehicle selfishly selects the shortest-time path, then this may not always

result in a global minimization of trip times. This result, due to Beckmann et al. [6], may be used as an argument against navigation systems that implement traffic guided routing. On the other hand, no proof exists that a sub-optimal solution is worse than the absence of any control.

The Comprehensive Automobile Traffic Control System (CACS), developed in Japan in the seventies, laid the foundations for the use of vehicular traffic sensors. The CACS system was designed to test the effectiveness of the provision of real-time traffic information to vehicles. Tests were run on a 4 km \times 7 km urban area in Tokyo, area which contained 85 intersections. Traffic information was recorded with 103 roadside units, 255 loop antennas and 1,000 taxis. The impact of traffic information dissemination was observed on 330 CACS vehicles. After this seminal work, many researchers have begun to investigate the impact of traffic information dissemination on traffic flows.

In [7] authors implement an integrated traffic flow, behavioral and traffic information model. The goal of this work is to show the impact of traffic information on a vehicular network, as the penetration ratio increases. Interestingly, in this study higher penetration ratios lead to poor overall performance. In [7] a fully informed traffic network attains the same performance of a system with no information feedback. However, this and other studies [8], [9] focus on very simple traffic networks, far from the complexity of the network that is under study in this paper.

In [10] a framework is defined to disseminate traffic information. This work describes in detail the architecture of a distributed traffic information system. However, authors do not investigate the effects the traffic information system may have on traffic. The evaluation only considers a 15 km straight street with 4 lanes in both directions. Similar observations were also independently found in [11].

More recently, authors of [12] implement a bi-directionally coupled simulator, integrating a vehicular and a telecommunications simulator. The analysis of the impact of a smart navigation system is limited to a test case with 200 vehicles that leave a single location and all head to the same destination.

A key role in this area is played by traffic measurement methodologies and by the technologies employed to gather the measurements [13]. One of the oldest and most widely spread technologies is induction loops. Induction loops are usually placed in the asphalt and provide punctual measurements for speed and traffic flow for that location. This metric suffers from a number of problems that limit its reliability. Intuitively, induction loops record speed information at certain locations on a street, thus they may result misleading with urban stop-and-go traffic. A more complete analysis of induction loops may be found in [14]. Video cameras are slowly replacing induction loops, but their widespread deployment is limited by their cost. The advantage in using video cameras is of recording end-to-end times rather than a punctual speed samples. While these methods are well established and their impact on traffic is well understood, we aim to compare them with a fully decentralized crowd-sourced solution.

Finally, in this paper we do not deal directly with security and privacy concerns. In vehicular networks, it is important

that the disseminated information is trustworthy, as it can affect driving decisions. Furthermore, the disseminated information can even raise safety issues (e.g., cause accidents) and raise privacy concerns. Privacy mechanisms could be built over our system to make sure that such a framework can be widely used [15], [16], [17]. Numerous trust mechanisms were also devised [18], [19], [20] to improve cooperation and quality of the disseminated information. Finally, security mechanisms [21], [22] can also be enforced to ensure safety and privacy. All these systems are orthogonal to our approach.

III. SCENARIO AND SYSTEM REQUIREMENTS

Vehicles provided with a computerized system can be aided in navigation by continuously assessing and correcting the best route prediction to a destination. To do this each vehicle should be capable of: i) sensing traffic information; ii) sharing it with neighboring vehicles (in an Ad-Hoc manner); iii) dynamically re-computing the best route to destination from the current position based on the collected information. Therefore, a Navigation System (NavSys) becomes an element of a distributed system that cooperatively collects and exchanges traffic conditions and, at the same time, a sophisticated traffic estimator, based on real-time information.

Such a system would be useful in many traffic scenarios, beyond regular traffic management operations. For instance, a VANET-based ITS may be used where infrastructure-based traffic monitoring systems are not deployable for various reasons. This is in fact quite common, only a few cities around the world implement an infrastructure capable of recording traffic data and of providing it to drivers. Other possible applications span from traffic management under uncommon conditions (i.e., accidents, disasters, evacuations) to less critical applications as constrained navigation (e.g., reach the closest station before gas runs out).

A number of basic characteristics identify a system able to offer the solution just described:

- 1) *Traffic Sensing*. Each vehicle should be able to act as a traffic sensor. The requirement for a traffic-sensing module is to measure a quantity closely related to traffic. Speed, traffic volume, traffic density and trip time are the most commonly used metrics;
- 2) *Traffic Information Dissemination*. Afterwards, vehicles will exchange the sampled information in an Ad-Hoc manner. The right traffic information should reach the right vehicles with the minimum delay. Excessive redundancy risks to congest the feedback channel, too little can lead to uninformed decisions;
- 3) *Traffic Estimation*. Finally, each vehicle should be able to independently evaluate the traffic conditions based on the traffic samples received through the network. Raw samples should be filtered and transformed in correct traffic estimates. For instance, a red traffic light phase should not be mistaken with a congested state. Moreover, as samples arrive at different rates, we need a model to estimate the traffic conditions when only a few or only very old samples are available.

To understand how such a system affects *global traffic conditions* it is important to evaluate the consequences of each

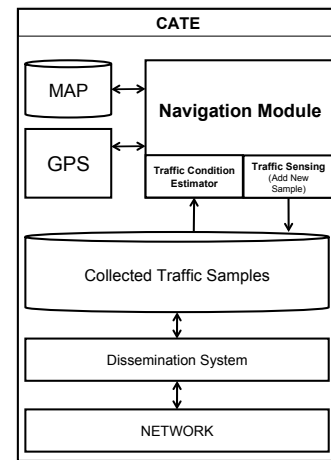


Fig. 1. CATE's architecture.

design choice. However, the complexity and the scale of the real system makes an “on the field” evaluation prohibitive. Therefore, there is a need for tools capable of producing realistic traffic data and handling dynamic routing. This is a very important piece of the puzzle. As we have seen in the related work section, in the past many studies [7], [8], [9] have concluded that informed navigation can lead to traffic disruption. In the following section, we will describe our prototype system and we will later describe the tools that we designed to study its impact on traffic.

IV. CATE: COMPUTER ASSISTED TRAVELING ENVIRONMENT

We now briefly examine the Computer Assisted Traveling Environment's (CATE) architecture (Figure 1). CATE is a smart navigation system designed to answer the requirements described the previous section. We will now provide more details about each one of the three key modules.

A. Traffic Sensing Module

CATE assumes that every vehicle can become a traffic sensor. Therefore, the vehicle's navigation systems should be able to accurately take these measurements. Street section delay is widely accepted as one of the most effective measures of the degree of congestion. Therefore, the vehicular navigation problem can be modeled as a search of the shortest path on a weighted graph. In such model, street sections are links, intersections are nodes, and a link's weight is given by the time required to traverse it. The graph model well applies to the majority of map databases [23] and navigation systems.

In CATE every time a vehicle exits a road segment, it creates a traffic sample of type $\{\text{linkID}, \text{delay}, \text{timeStamp}, \text{carID}\}$. The `linkID` is unique per street segment and direction throughout the vehicular network. The `delay` field represents the time spent by the vehicle on link `linkID`. The GPS and the map information is used to identify the time when the vehicle entered/exited a certain road segment. The `timeStamp` is the GPS time at which the sample was collected and more specifically the time when the vehicle exited the measured road segment.

B. Dissemination Module

The dissemination module deals primarily with spreading these samples throughout the vehicular Ad-Hoc network in an efficient way.

Previous studies [24], [25], [26], [27], [28] show how the performance of traditional Ad-Hoc routing dissemination is heavily affected by vehicle density in urban areas. Conversely, gossip-based Ad-Hoc routing is very effective [29] in disseminating large amounts of information in dense networks and this is why we choose a utility-based gossip model in CATE. We will now provide more details about our traffic sample dissemination protocol.

Periodically, CATE selects a subset of the samples that are available in the buffer (notice that the buffer contains both the vehicle's collected information and information generated by other vehicles) and broadcasts this information to its neighbors. One-hop neighbors will combine the received samples with those that were already in their buffer, and later spread them even further.

A key element of the dissemination module is the *sample selection algorithm*: how to select which subset of information to broadcast, assuming only a fraction of a vehicle's knowledge can be sent within given bandwidth restrictions. Therefore, we designed a simple mechanism to *prioritize the samples*. To take these decisions each sample is ranked with the help of *utility function*: a metric that represents the effectiveness of each sample. Afterwards the K links with the highest utility are broadcasted to the neighbors.

There are numerous ways to design an appropriate utility function but we chose to emphasize on two simple factors:

- **Prefer fresh samples:** rank the links based on how recent is the information. In that case $r_{mostRecent}$ is calculated based on the $timeStamp_{linkID}$.
- **Maximize coverage:** rank the links based on how rare is the disseminated information. In that case $r_{lessbroadcasted}$ is defined as $1/NB_{LINK}$ where NB_{LINK} is the number of times information about this *link* was received during the last t minutes (if $NB_{LINK} = 0$ we move the link to the top of the ranking).

Afterwards, we linearly combine these two rankings to prioritize fresh links but also to encourage high geographical coverage of the disseminated information. Therefore, our utility is $U_{Link} = r_{mostRecent} + r_{lessbroadcasted}$. As we shall see in Section V, this simple approach provides satisfactory results in a 4 km x 7 km realistic urban scenario. Numerous optimizations can improve the information dissemination procedure (e.g., use a utility that includes topology information to disseminate only information that is needed by nearby vehicles). However, an in depth study of such algorithms is beyond the scope of this paper.

C. Traffic Estimation and Dynamic Routing Module

As we described before, each vehicle will individually collect a set of measurements (samples) through the network interface. The key challenge is how to estimate the current traffic conditions and dynamically re-route the vehicle.

CATE periodically computes the route of a vehicle using the modified Dijkstra algorithm described in [30] (although other algorithms specifically devised for vehicular networks can be easily integrated [31]). In practice, the modified version of Dijkstra only considers those links (road segments) that are inside the geographical area between the current position and the destination. Each link has a weight w that allows Dijkstra to find the shortest path. In our case, *the weight represents the updated traffic condition of the segment* (i.e., how long a vehicle is expected to be on the road segment based on the information that it managed to collect via the network).

The problem of estimating a link's traffic condition based on the collected samples is not trivial due to the following reasons:

First of all, *noise in the observations*: although traffic conditions do not change rapidly over time (traffic jam's dynamics are relatively slow), sample delays can be quite noisy. Distinct vehicles may drive through the same street segment at the same time, but at a different pace. Some vehicles may stop at a traffic light, at a pedestrian crossing or to pick up a passenger, while other vehicles rush through the street segment. The result is that samples may vary significantly, although collected close in time.

Secondly, *sample rate variation*: there is no guarantee that a link's traffic samples will be received at a fixed rate (it depends on traffic and on the dissemination strategy). For example, a vehicle may receive multiple old samples and just a few fresh. One key problem is how to weight this information (age of samples) in order to calculate the best possible estimation of the current conditions.

Finally, *absence of information*: no recent information may be received for some road sections. The two possible choices are to use weights that represent either historical data (e.g., typical average speed at that time on that road segment) or speed limits. However, we need to identify how quickly the information is considered as obsolete.

To answer these questions, we tested a number of solutions in order to interpret the collected samples to provide an accurate estimation of the current traffic conditions:

- **Default:** a link's weight is computed dividing its length by its speed limit:

$$W_{linkID} = \frac{LinkLength}{SpeedLimit} \quad (1)$$

CATE sets this as the default weight when no information is known about a link.

- **Most Recent Estimate:** for each link CATE selects the most recent sample in terms of `timeStamp` (not the time this vehicle received it):

$$W_{linkID} = Delay_{linkID, mostRecentSample} \quad (2)$$

This is the simplest algorithm as it discards all the samples that were collected earlier. However, a link's weight can heavily fluctuate, since samples can vary rapidly. Surprisingly, the results illustrated in Section V show that this approach achieves a satisfactory performance. From a theoretical viewpoint, this is justified by observing that end-to-end times on road segments follow a bimodal

distribution [32]. If the distribution that models the low congested state and the distribution that models the highly congested state are not very spread, considering that it takes some time to transit from one state to the other, any recent traversal time sample can correctly represent the current state of the link.

- *Bayes Estimate*: We use a simple Bayesian estimator to predict the current traffic conditions by using a large number of samples taken at different times:

$$W_{linkID} = (1 - \alpha) * Delay_{NewSample} + \alpha * CurrentWeight_{LinkID} \quad (3)$$

where α is a parameter tuned by the sample's age (so that older samples do not excessively weight).

- *Bayes with Aging Estimate*: The same as before, but the absence of information moves the weight back to the default value given by the free flow traversal time:

$$W_{linkID} = (1 - c) * BayesWeight_{linkID} + c * DefaultWeight_{LinkID} \quad (4)$$

c is an aging factor computed as:

$$c = \frac{Min(curTime - recentSampleTime, maxAge)}{maxAge} \quad (5)$$

This approach differs from [10], where information is always assumed to be known, in interpreting no information as no congestion, and, therefore, accordingly dropping link weights.

V. THE EVALUATION PLATFORM

To evaluate CATE, we designed and implemented a tool that simulates both dynamic vehicular navigation (mobility) and the VANET dissemination. Solutions that couple a mobility simulator and a telecommunications network simulator can be found in [12], [33], [34], [35]. Differently from in previous work, we integrate QualNet [36], a communications network simulator specifically designed for wireless networks, and MobiDense [37], a mobility simulator we designed and implemented. These two simulators constantly interact: *future mobility decisions are influenced by the network dissemination (e.g., collected information), and the network dissemination is influenced by the mobility patterns (location of the vehicles)*. A depiction of the interactions between the two tools is shown in Figure 2.

A. MobiDense

MobiDense [37] is a mobility simulator which combines topology and traffic flow information to generate a mobility trace. We chose MobiDense as it allows us to dynamically modify the road weights and because it can directly plug in to QualNet. It is able to simulate vehicular mobility on real road-network topologies that are extracted from digital maps. Moreover, we use detailed information about the type




Fig. 2. *MobiDense and QualNet Interactions.*

of the street segments (i.e., two-way or one-way roads, speed limits, number of lanes), the probability of stopping at an intersection, traffic lights locations, traffic light timings and phases, etc. When no additional information is provided with the digital map, the street segment capacities are used to estimate the enforcing speed limits, and to adapt intersection stop probabilities and red/green time phases of traffic lights. Consequently, queues may build up at intersections and propagate backwards. Besides topological information, MobiDense also requires a traffic flow model: the starting point and time, and the destination of all vehicles to be simulated. Intermediate routing decisions are taken individually by each vehicle's routing module based on the topology and the street weights that are used.

In terms of interaction with QualNet, at each time step t for each vehicle v that is now at location $v < x_t, y_t >$ MobiDense calculates its position for the next time step $v < x_{t+1}, y_{t+1} >$. To perform this task, MobiDense simulates v 's mobility based on the route that it selected to reach its final destination, the road topology and the other vehicles in the simulation (e.g., traffic queues, speed limits, traffic lights, car-following models, etc). However, periodically MobiDense re-calculates v 's route to its final destination based on the traffic samples collected via the network (the samples collected per vehicle are placed in the "Collected Traffic Information" database that shown in Figure 2). Furthermore, since each vehicle acts as a traffic sensor, MobiDense also adds samples to the database: when a vehicle exits a road segment a sample for this link is added in the database. Later this sample will be disseminated via the network (initially to v 's neighbors and later further away).

In conclusion, MobiDense takes the role of a) a smart navigation system that can dynamically route each vehicle based on the estimated traffic conditions that were received via the network, b) a traffic sample sensor and c) a mobility simulator that moves the vehicles based on the topology and the real traffic.

B. QualNet

The wireless network dissemination modules are implemented in QualNet [36], a well-known network simulator particularly suited for the simulation of wireless networks.

First of all, at each instance t the vehicles are placed at the locations that were instructed by MobiDense. Since we are using short-range radio these locations will determine connectivity.

To disseminate information, each vehicle periodically broadcasts samples that are found in “Collected Traffic Information” database (a shared buffer between the two simulators). As described in the previous sections, these broadcasts contain only a subset of the samples that are in the database. Similarly, when a vehicle receives a message from the network, it adds the newly received samples to the database. MobiDense will later use them to dynamically re-route the vehicle and QualNet will further disseminate them.

Notice that QualNet and MobiDense continually exchange information (they synchronize once every second). MobiDense provides the vehicles’ positions and the streets’ traversal times to QualNet while QualNet is used to gossip the information.

VI. STUDYING THE IMPACT OF VANET DISSEMINATION ON TRAFFIC

We will now move to our main contribution: the evaluation of the impact of the system presented in Section IV. The goal of this work is to establish whether a distributed (Ad-Hoc) system such as CATE, when used by all vehicles, can reduce traffic congestion. The main performance metric is the vehicle’s total trip time, but also other aspects should be considered. First, we will study how performance improvements, if any, are distributed among vehicles. Second, we will measure how quickly the information is disseminated. Finally, we will quantify the amount of communication traffic that is produced.

For our evaluation, we use a scenario based on downtown area of Portland, Oregon. The area is approximately 4 km x 7 km, a map is shown in Figure 5. The area includes 4,968 streets and 3,429 intersections. About 16,500 vehicles are involved in the simulation. Each vehicle’s journey information (start time and location and end point) is extracted from large-scale surveys on the population of Portland, Oregon. This information is extremely valuable as it is very important to have realistic traffic flows that follow a certain spatio-temporal distribution. Origin-destination pairs are then plugged into our coupled mobility-telecommunication simulator.

We use the same source to extract information such as traffic light positions and delays, intersection stop probabilities, speed limits and road capacities. MobiDense produces traces, in absence of information dissemination between vehicles, which reproduce in terms of flows the original survey observations.

A. Traffic Information Evaluation

Each vehicle stores traffic samples, grouped by `linkID`, in a local buffer. In case no information is known about a link, all the strategies we implement assume there is no traffic on such link (we assume that vehicles traverse the link at free flow speed). We now compare the three different traffic estimation algorithms described in Section IV: (a) *Most Recent Estimate*; (b) *Bayes*; (c) *Bayes with Aging*. We additionally compare these strategies to the case where no information is disseminated.

Traffic flows are adjusted from 33% of the traffic that is found in the original survey to double the actual flow values. For example, if an average of 100 vehicles per hour enter

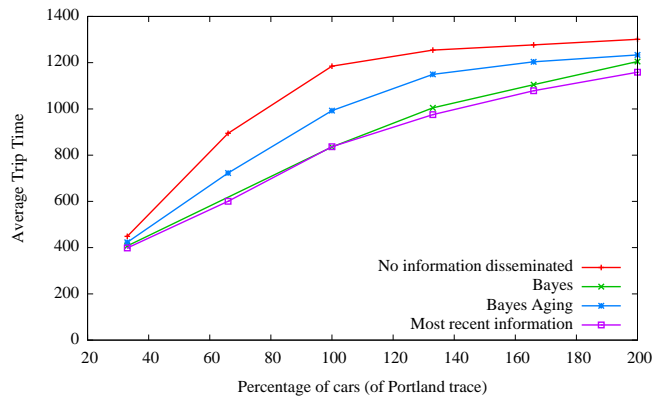


Fig. 3. Trip times for different information handling strategies (in seconds).

the map from a certain intersection in the original survey, we begin the simulation with a flow of 33 vehicles per hour. Therefore, we simulate the network observing its behavior with low densities and we then gradually increase traffic to reach typical morning traffic volumes.

As we see in Figure 3, in absence of information feedback, when density increases overall trip times quickly rise from 400 seconds, about 7 minutes, to 1200 seconds, about 20 minutes, on average. When CATE is used, trip times drop. A higher absolute improvement is observed under normal traffic congestion but the gain is surprisingly less when we have higher than normal traffic conditions. This result radically differs from what is found in [7], [8], [9], [38]. An intuitive explanation may be found by observing Figure 5. As we can see in Figure 5.a, traffic is mainly localized on the freeways and on the bridges that traverse the river. However, traffic is not all generated by vehicles that traverse the river or that take a freeway. Many vehicles could reach their destinations through alternate routes, but they do not. In Figure 5.b we see the result of using CATE. Many more yellow links (i.e., slightly congested links) and fewer red links (i.e., heavily congested) are present. In fact, heavily congested links are substituted by a number of slightly congested (yellow) links in Figure 5.b. This indicates that when vehicles collected the traffic information, they were diverted from heavily congested areas to areas that were previously not congested, creating some minor congestion.

The space for improvement is high because the topology we are here analyzing is realistic and more than one path is usually available to reach a location. At the center of the map, in the Manhattan-like section, we can see an increase of yellow links when using CATE. Under normal morning traffic conditions, the average trip time is 29% lower, which is a significant improvement.

In terms of weight calculation algorithms, we observe that the most recent information and Bayes strategies provide the best results in this simulated scenario. This happens because if there is traffic on a link, end-to-end times increase considerably, while they otherwise oscillate slightly above the free flow delay time. Bayes with aging still improves, but not as much as the other two methods since traffic conditions are not rapidly changing.

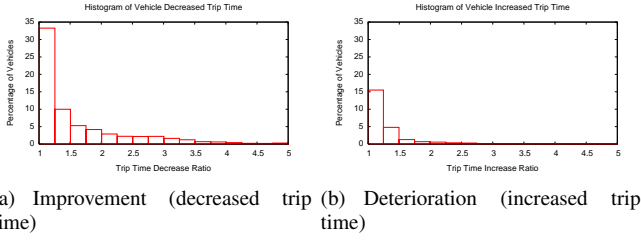


Fig. 4. Histogram of trip-times loss/gain.




Fig. 5. Map of speed [best viewed in color]. Green streets are not congested. Yellow areas show average speed slightly lower than speed limit. In red streets segments the average speed is much lower than speed limit.

It is important to understand how travel time improvements are distributed. Figure 4 presents a histogram of the trip times gain/loss for normal traffic conditions when dissemination is used. More specifically, gain ratio (Figure 4(a)) is defined as $ratio = \frac{oldtime}{newtime}$. For example a ratio of 2 means that the vehicle halved its trip time, a ratio of 3 that it needed one third, etc. Similarly, deterioration time (Figure 4(b)) is defined as $ratio = \frac{newtime}{oldtime}$ (a ratio of 2 means the vehicle doubled its trip time). As we observe, a large number of vehicles, 34%, saved 20% of the time (ratio 1.25 means new time is 1/1.25 of the old time). There were also luckier vehicles able to avoid big traffic queues and complete their journey two or three times faster (ratio 2 and 3). In total 64% of the vehicles saved time. However, at the same time, we see that some of the drivers required more time when our application was used. This is due to some of the traffic being diverted into smaller roads that, consequently, become busier. However, we can observe that far fewer drivers have their time increased rather than decreased and their trip times are no more than two times longer. Finally, 23% of vehicles were not really affected ($\pm 10\%$ trip time).

Finally, in terms of penetration ratio, we experimented with a range of 10% to 100%. As expected, when less vehicles are equipped with CATE, the system's performance is slightly worse. However, in our experiments we noticed that when only 34% of the vehicles are equipped with CATE the performance is comparable to higher penetration ratios.

B. Traffic Information Dissemination Quality

We also want to understand how well CATE disseminates traffic information, giving a close representation of real traffic conditions to each vehicle. In fact, the results we presented in Figure 3 may derive from an unfair dissemination of information. We should remember that, in simple scenarios [7], [8], [9], [38], it has been shown that a fully informed traffic network can deteriorate traffic performance. Therefore, random inconsistencies in traffic information dissemination may be inducing a better traffic behavior.

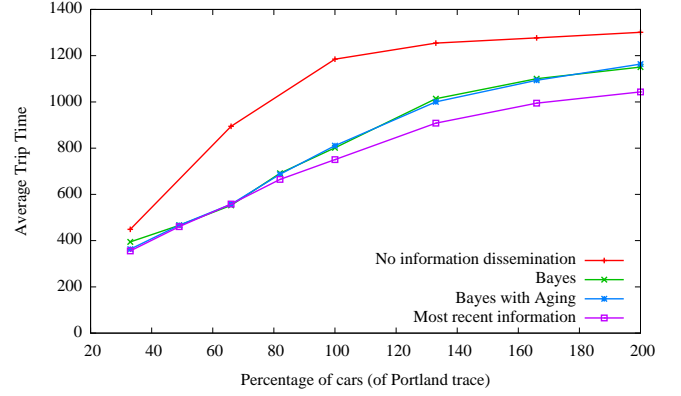


Fig. 6. Trip times (seconds) when full knowledge is available instantly to all the vehicles (best information dissemination case).

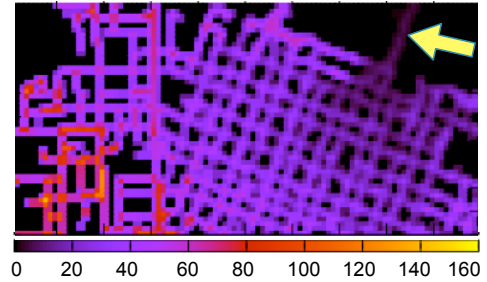


Fig. 7. 2D-Heatmap of age of received information (in seconds) about the link highlighted by the arrow (bridge). Vehicles away from the bridge receive older traffic information. [best viewed in color].

To estimate the impact of the dissemination protocol on the system's performance, we compute the overall average travel time when all traffic information is immediately available at all vehicles. This is the infinite bandwidth/zero delay scenario, a full-knowledge scenario where all vehicles know all possible information in real-time. We then observe the variation in aggregated average traffic trip time, between a fleet of vehicles that gossip information and a fleet of vehicles that immediately receive all the available information. Please notice that this mode would not be possible in reality and we just use it to compare with full-knowledge scenario.

Figure 6 shows the results for the same traffic estimation methods used before, yet now information is not collected with the dissemination protocol (i.e., gossip). All the collected information is instantly available when a vehicle requires running its routing algorithm. We observe that the same trends appear as when CATE is used. The comparison with Figure 3 reveals that trip times are slightly smaller. This result is very interesting as we can conclude that informed vehicles can reduce the overall average trip time and that CATE is performing well, giving an updated picture of the network to each vehicle.

To better understand how recent information is received at a vehicle, we analyze the information propagation speed on the map. Figure 7 shows a zoomed area of the map (near the central section, between Burnside Bridge and Morrison Bridge). In this graph, we plot the average age of information about the bridge pointed by the arrow, Morrison Bridge. In

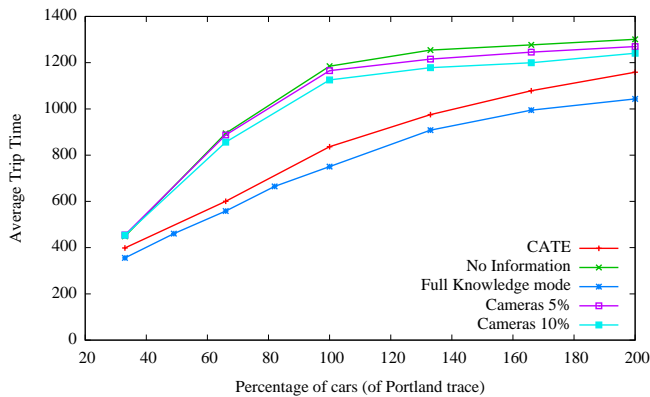


Fig. 8. Infrastructure versus CATE: average trip time (seconds).

nearby areas this information is on average less than one minute old. In areas that are about two kilometers away, information is on average about three minutes old. In fact, in the whole 4 km x 7 km simulation area we could rarely find vehicles that were using information older than fifteen minutes. This explains why the results shown in Figure 6 and 3 are so close: CATE performs close to full-knowledge since traffic trends (i.e., as congestion build up) are slower than information dissemination, thus giving vehicles enough time to react.

C. Infrastructure vs. Infrastructure-less Probing

We here compare CATE to state of the art solutions that monitor selected street segments using video cameras or induction loops, and disseminate traffic information using cellular networks (e.g. 3G) or FM stations. With such systems, all vehicles can instantly access the same updated and accurate information about the instrumented roads. In CATE we have the opposite situation, information is collected by all the vehicles (and, thus, on almost all street segments), but information is not as recent due to dissemination delay. Additionally, with CATE, distinct vehicles are likely to have a different perspective of the traffic situation.

For our comparison, we monitor 5% and 10% of the streets of downtown Portland. These streets are not chosen randomly, we chose the most congested 5%-10%. Therefore, the delay information about the most congested streets is instantly delivered to all the vehicles.

Figure 8 shows the total average trip times as traffic flows increase. The infrastructure-based solution is outperformed due to its lack of flexibility. When a vehicle learns that a monitored street is congested, it clearly avoids it. However, as a result vehicles may congest streets with no traffic monitoring capabilities. On the other hand, CATE collects information about all streets using each vehicle as a mobile sensor. Therefore, the build up of new congestion points is reported. Information might be delayed, but it is recent enough to avoid congestion hotspots.

D. Studying the Impact of Misbehaving Nodes

Furthermore, although trust and security protocols can be used to ensure that the vehicles cooperate and are accountable

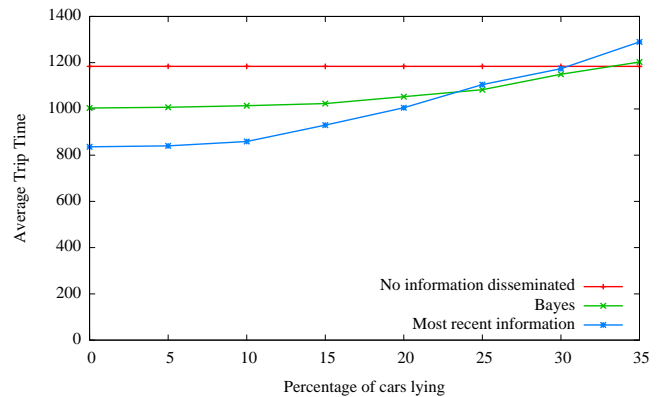


Fig. 9. Performance when a percentage of nodes are misbehaving.

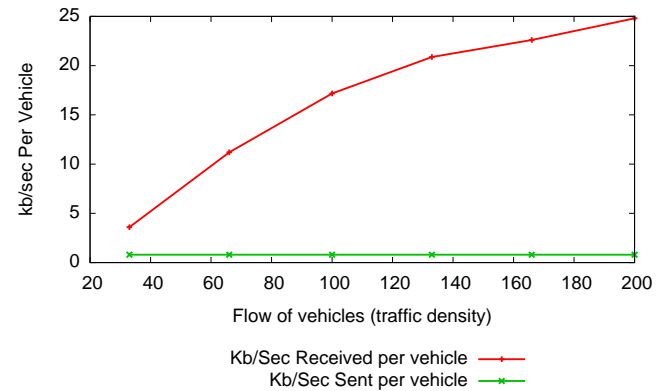


Fig. 10. Network data rates.

for the information that they share (e.g., [18], [19], [20], [21], [22]), we were also interested in identifying what is the impact of misbehaving nodes on the traffic conditions.

In our scenario, misbehaving nodes are nodes that intentionally spread wrong information about a pre-agreed selection of road segments. More specifically, in our simulation we would like to identify the worse case scenario: all the misbehaving nodes will advertise that all the bridges are fully congested apart from one: I405 (the leftmost bridge in Figure 5). If a large percentage of nodes spread this information, this will force a large number of vehicles to use highway I405.

In Figure 9 we plot how different numbers of misbehaving nodes affect trip times (0% is the reference point). Clearly, a small percentage (< 10%) does not affect our system. As the percentage of malicious nodes increases to a significant size (> 22%), we observe that Bayes achieves better performance than using the most recent estimate. This happens because Bayes is averaging multiple samples in order to create a more accurate estimation. In any case, we observe that it requires an unrealistically large amount of misbehaving nodes (> 30%) to impact worse driving times than the scenarios that do not use our system.

E. Gossip Network Overhead

We measured the dissemination protocol network overhead in order to estimate the communication's network congestion.

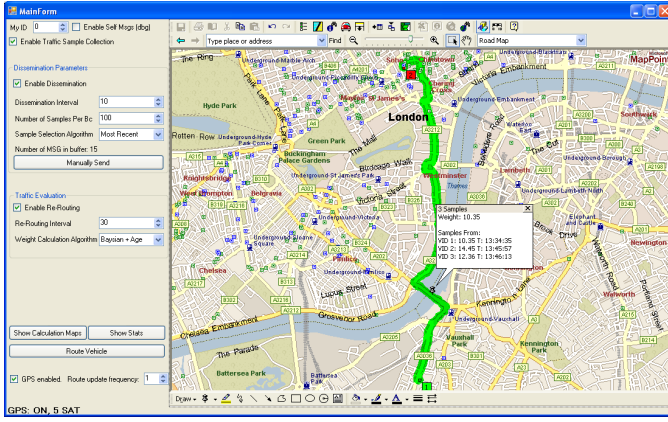


Fig. 11. Our Implementation Prototype.

In all simulations, we used a 10 seconds gossip interval and a 2000 bytes sending buffer. Figure 10 presents the transfer rates experienced by each node. As it can be observed, plenty of space is available for more data. Vehicles only receive data at 16kbps during normal morning traffic conditions, while devices compliant to the 802.11p [39] standard are expected to provide transmission rates in the order of tens of megabits.

VII. ON THE ROAD

Finally, we were interested in examining whether such a system is feasible. There are many technical challenges in building a system like CATE. First of all, vehicles should be able to accurately identify and collect samples using their GPS and maps. Furthermore, there are computationally demanding tasks (e.g., traffic estimation based on received samples, re-routing, etc). Finally, real performance problems such as radio range and radio propagation speed can be evaluated. Therefore, we implemented a fully working prototype of CATE. This prototype allows us to further understand the challenges in such a system.

Figure 11 shows our implementation’s graphical user interface coded in C# and using Microsoft MapPoint as our navigation system API. The prototype, thought for experimentation, enables a user to manually tune parameters such as the route re-computation interval, the dissemination interval, the sample selection and the weight calculation algorithms.

We here assess the practicality of deploying CATE implementing its main components and performing a set of connectivity experiments. Unfortunately, a more interesting set of experiments that involve traffic information gathering and dissemination and vehicular navigation is unfeasible.

We use the UCLA Campus Vehicular Testbed (C-VeT) [40]. C-VeT offers both vehicle-to-vehicle and vehicle-to-infrastructure connectivity, a virtualized shared environment, and a number of tools to record a system’s behavior and performance.

A. Experimental setup

During the experiments, each vehicle carries a PC equipped with a Zyxel AG-225H card (IEEE 802.11a/b/g compliant), and a SIRF STAR III based GPS receiver. The hosts were

assigned a static IP as well as a predefined SSID. A high gain omnidirectional antenna was installed on each vehicle’s rooftop, and the line of sight range was approximately 250 meters.

We designed two classes of experiments to study the feasibility of CATE. First, we performed connectivity experiments to evaluate the amount of traffic samples that can be transmitted between two vehicles that are directed in opposite directions. Second, we tested CATE’s basic communication features, namely, its neighborhood discovery scheme and its dissemination scheme. Eight vehicles were involved in this second test.

B. Experiments

We performed a number of experiments using our prototype implementation, aiming to examine the feasibility of such a system.

1) *Connectivity*: in order to assess the connectivity between two nodes surrounded by traffic we performed two experiments. First, we measured the throughput achieved between two static vehicles. In this way, we define the best case. Second, we measured how much traffic samples could be disseminated between two cars traveling in opposite directions at 30 miles per hour (or 13.4 meters per second). Each experiment was repeated 10 times.

In the static case, at the UDP layer the average throughput is 30.004 Mbps, the peak throughput is 30.3 Mbps, and the minimum throughput 28 Mbps. Two vehicles were driven following the normal traffic, and never exceeded the speed of 30 miles per hour. Once they were in radio range a stream of UDP packets flowed from the first vehicle to the second. We measured the amount of data transferred during each contact. On average, β received 7.24 MB during an average contact time of 15 seconds.

2) *Dissemination*: CATE was deployed on C-VeT, and tested in its basic features. In fact, we tested the neighborhood discovery protocol and the dissemination protocol. Experiments were performed using eight vehicles moving in UCLA’s campus.

The neighborhood discovery protocol transmitted hello packets at 1 Mbps with no acknowledgements, while gossip packets were transmitted at the nominal rate of 54 Mbps.

We performed two tests:

One-hop gossiping: In this experiment, each vehicle gossiped information to all other vehicles in range. We measured the average amount of traffic samples transferred during a contact among vehicles; on average, 2.35 MB were transferred (about 125000 traffic samples). This value is lower than the experiment with two cars because: (a) the channel load is higher (each car is transmitting to 2-3 neighbors on average; (b) the mobility pattern created a long tail distribution of contact times ranging from three seconds to thirty seconds. We observed a peak data transfer of 3.35 MB and a minimum value of 1.76 MB.

2 MB traffic-sample dissemination: In the second test, only one car behaved as a source of a 2 MB traffic-update chunk that represents a full map update. We here aimed at estimating

the average delay experienced when transferring a big amount of traffic samples from this vehicle to all the others. In this experiment, we used CATE's dissemination scheme, setting the dissemination interval to 1 second. On average the 2 MB samples reached all other vehicles in 125 seconds. On average, the first vehicle received the entire amount of data in less than 20 seconds, the majority of the other vehicles (5 out of 8) received the data in 72.5 seconds, and the last two vehicles received the data in 118 and 125 seconds respectively.

While the experiments described in this section are far from realistic since they do not consider vehicular re-routing based on traffic estimations, the experiments show the feasibility of the proposed application. In fact, CATE is easy to deploy with off the shelf equipment and the computation power available in any navigation system or PDA device. Moreover, the bandwidth provided by available wireless technologies is more than sufficient to support the low data rates that an Ad-Hoc ITS requires.

VIII. CONCLUSION

In this paper, we quantified the effects of deploying a decentralized traffic based navigation system in downtown Portland, Oregon. Moreover, we analyzed the effects full knowledge and partial monitoring would have on traffic congestion. First of all, our results show that a decentralized approach can reduce traffic congestion in a realistic scenario. This result is not trivial, since there are a number of previous works that point in the opposite direction. Second, our results show that a gossip scheme performs very well, both in terms of vehicular traffic and of telecommunications traffic overhead. In fact, when nodes are provided with full information about traffic with no delay, the average travel time only slightly decreases. Third, we show that monitoring only a subset of streets, even when these are the most congested streets, can lead to unsatisfactory results. In terms of future work, we are interested to study the joint optimization of information dissemination and vehicle's trip times. Moreover, smart distribution techniques of traffic information could avoid the build up of secondary congestion generated by the dissemination of identical information (e.g., avoiding a closed road can congest the obvious second best road). In conclusion, smart navigation represents a powerful and cost efficient tool, which together with others (e.g., use of public transportation, etc.), can combat the increase of traffic congestion in urban areas.

Acknowledgments: This work was partially supported by US Army under the DURIP program, the National Science Foundation under the BBN-GENI program Grant #1797, Cisco fund #2010-03937, the Macau Polytechnic Research Fund grant #RP/ESAP-03/2010 and the Italian FIRB Project DAMASCO.

REFERENCES

- [1] F. Calabrese, M. Colonna, P. Lovisolo, D. Parata and C. Ratti, "Real-Time Urban Monitoring Using Cell Phones: A Case Study in Rome," *IEEE Trans. Intelligent Transportation Systems*, vol. 12, pp. 141-151, Mar. 2011.
- [2] R.G. Herrtwich, "Communicating Vehicles - Communicating Roadways: New Approaches to Driver Information and Road Safety," *Invited Talk at the 11th Annual ACM International Conference on Mobile Computing and Networking*, 2005.
- [3] R. Resendes, "The New Grand Challenge - Deploying Vehicle Communications," *Keynote Address at the 5th ACM International Workshop on Vehicular Inter-NETworking*, 2008.
- [4] Y. Yang and R. Bagrodia, "Evaluation of VANET-based Advanced Intelligent Transportation Systems," in *Proc. of the 6th ACM International Workshop on Vehicular Inter-NETworking*, 2009, pp. 3-12.
- [5] S. Eubank, H. Guclu, V. Kumar, M. Marathe, A. Srinivasan, Z. Toroczkai, N. Wang, "Modelling disease outbreaks in realistic urban social networks," *Nature*, vol. 429, pp. 180-184., Nov. 2004.
- [6] M. Beckmann, C.B. McGuire and B. Winsten, "Studies in the Economics of Transportation," Yale University Press, Research Memoranda RM-1488-PR, 1955.
- [7] R. Jayakrishnan, H. Mahmassani and T.Y. Hu, "An evaluation tool for advanced traffic information and management systems in urban networks," *Elsevier Transportation Research Part C: Emerging Technologies*, vol. 2, pp. 126-147, Sep. 1994.
- [8] R. Arnott, A. de Palma and R. Linsley, "Does providing information to drivers reduce traffic congestion?," *Elsevier Transportation Research A: General*, vol. 25, pp. 309-318, Sep. 1991.
- [9] Al-Deek, H.M., Khattak, A.J., Thananjeyan, P., "A combined traveler behavior and system performance model with advanced traveler information systems," *Elsevier Transportation Research Part A: Policy and Practice*, vol. 32, pp. 479-493, Sep. 1998.
- [10] T. Nadeem, S. Dashtinezhad, C. Liao and L. Iftode, "TrafficView: traffic data dissemination using car-to-car communication," *ACM SIGMOBILE Mobile Computing and Communication Review*, vol. 8, pp. 6-19, Jul. 2004.
- [11] L. Wischhof, A. Ebner and H. Rohling, "Information Dissemination in Self-Organizing Intervehicle Networks," *IEEE Trans. Intelligent Transportation Systems*, vol. 6, pp. 90-101, Mar. 2005.
- [12] C. Sommer, R. German and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Trans. Mobile Computing*, vol. 10, pp. 3-15, Jan. 2011.
- [13] F. Soriguera and F. Robuste, "Requiem for Freeway Travel Time Estimation Methods Based on Blind Speed Interpolations Between Point Measurements," *IEEE Trans. Intelligent Transportation Systems*, vol. 12, pp. 291-297, Mar. 2011.
- [14] M. Westerman, R. Litjens and J. Linnartz, "Integration Of Probe Vehicle And Induction Loop Data: Estimation Of Travel Times And Automatic Incident Detection," University of California, Berkeley, Tech. Rep. UCB-ITS-PRR-96-13, 1996.
- [15] K. Sampiget, L. Huang, M. Li, R. Poovendran, K. Matsuura and K. Sezaki, "Caravan: Providing Location Privacy for Vanet," *ESCAR Embedded Security in Cars*, 2005.
- [16] M. Gerlach, "Assessing and Improving Privacy in VANETs," *ESCAR Embedded Security in Cars*, 2006.
- [17] M. Gerlach and F. Guttler, "Privacy in VANETs using Changing Pseudonyms - Ideal and Real," in *Proc. of the 65th Spring Vehicular Technology Conference*, 2007, pp. 2521-2525.
- [18] J. Serna, J. Luna and M. Medina, "Geolocation-Based Trust for Vanet's Privacy," in *Proc. of the International Symposium on Information Assurance and Security*, 2008, pp. 287-290.
- [19] M. Gerlach, "Trust for Vehicular Applications," in *Proc. of the 8th International Symposium on Autonomous Decentralized Systems*, 2007, pp. 295-304.
- [20] Z. Wang and C. Chigan, "Cooperation Enhancement for Message Transmission in VANETs," *Journal of Wireless Personal Communication*, vol. 43, pp. 141-156, 2007.
- [21] P. Papadim, L. Buttyan, T. Holczer, E. Schoch, J. Freudiger, M. Raya, Z. Ma, F. Kargl, A. Kung and J.P. Hubaux, "Secure Vehicular Communications: Design and Architecture," *IEEE Communications Magazine*, vol. 46, pp. 2-8, Nov. 2008.
- [22] J.J. Haas, Y.C. Hu and K.P. Laberteaux, "Design and Analysis of a Lightweight Certificate Revocation Mechanism for VANET," in *Proc. of the 6th ACM international workshop on Vehicular Inter-NETworking*, 2009, pp. 89-98.
- [23] <http://tiger.census.gov/>
- [24] F. Bai, N. Sadagopan and A. Helmy, "The important framework for analyzing the impact of mobility on performance of routing for ad hoc networks," *Elsevier Ad Hoc Networks Journal*, vol. 1, pp. 383-403, Nov. 2003.
- [25] G. Marfia, G. Pau, E. De Sena, E. Giordano and M. Gerla, "Evaluating Vehicle Network Strategies for Downtown Portland: Opportunistic Infrastructure and the Importance of Realistic Mobility Models," in *Proc. of the 1st MobiSys International Workshop on Mobile Opportunistic Networking*, 2007, pp. 47-51.

- [26] I. Leontiadis and C. Mascolo. "Opportunistic Spatio-Temporal Dissemination System for Vehicular Networks," in *Proc. of the 1st international MobiSys workshop on Mobile opportunistic networking*, 2007, pp. 39-46.
- [27] I. Leontiadis, P. Costa and C. Mascolo, "A hybrid approach for content-based publish/subscribe in vehicular networks," *Elsevier Pervasive and Mobile Computing*, vol. 5, pp. 697-713, Dec. 2009.
- [28] C.E. Palazzi, M. Roccetti and S. Ferretti, "An Inter-Vehicular Communication Architecture for Safety and Entertainment", *IEEE Trans. Intelligent Transportation Systems*, vol. 11, pp. 90-99, Mar. 2010.
- [29] Z.J. Haas, J.Y. Halpern and L. Li, "Gossip-based ad hoc routing," *IEEE/ACM Trans. Networking*, vol. 14, pp. 479-491, Jun. 2006.
- [30] R. Sedgewick and J.S. Vitter, "Shortest Paths In Euclidean Graphs", *Springer Algorithmica*, vol. 1, pp. 31-48, 1986.
- [31] S. Qing and W. Xiaofan, "Efficient Routing on Large Road Networks Using Hierarchical Communities," *IEEE Trans. Intelligent Transportation Systems*, vol. 12, pp. 132-140, Mar. 2011.
- [32] S. Tepley and G.D. Evans, "Evaluation of the Quality of Signal Progression by Delay Distributions," *TRB Transportation Research Record*, no. 1225, pp. 1-7, 1989.
- [33] C. Lochert, A. Barthels, A. Cervantes, M. Mauve and M. Caliskan, "Multiple simulator interlinking environment for IVC," in *Proc. of the 2nd ACM International Workshop on Vehicular Ad Hoc Networks*, 2005, pp. 87-88.
- [34] S.Y. Wang, C.L. Chou, Y.H. Chiu, Y.S. Tseng, M.S. Hsu, Y.W. Cheng, W.L. Liu, T.W. Ho, "An Integrated Simulation Platform for Vehicular Traffic Communication and Network Researches," in *Proc. of the 1st IEEE International Symposium on Wireless Vehicular Communications*, 2007, pp. 2081-2085.
- [35] Y. Pigne', G. Danoy and P. Bouvry, "A platform for realistic online vehicular network management," in *Proc. of the 2nd IEEE International Workshop on Management of Emerging Networks and Services*, 2010, pp. 615-619.
- [36] <http://www.scalable-networks.com/>
- [37] G. Marfia, P. Lutterotti and G. Pau, "MobiTools: An Integrated Toolchain for Mobile Ad Hoc Networks," University of California, Los Angeles, Tech. Rep. #070019, 2007.
- [38] K. Kobayashi, "Information and rational expectations and network equilibria - an analytical perspective for route guidance systems," *Springer The Annals of Regional Science*, vol. 28, pp. 369-393, Dec. 1994.
- [39] http://grouper.ieee.org/groups/802/11/Reports/tgp_update.htm
- [40] <http://vehicular.cs.ucla.edu/>



Dr. Ilias Leontiadis is currently a Research Associate in the Computer Laboratory, University of Cambridge. He holds a PhD in Computer Science from University College London, United Kingdom (2009) and a MSc in Computer Systems from the University of Ioannina, Greece (2004). He also spent some time at University of California Los Angeles (UCLA) and Microsoft Research Cambridge. His research interests include vehicular networks, delay tolerant networking, mobile networking and systems, wireless sensor systems and mobility modeling.

More information can be found at <http://www.cl.cam.ac.uk/~il235/>.



Dr. Gustavo Marfia received the Laurea degree in telecommunications engineering from the University of Pisa in Tuscany, Italy in 2003, and the Ph.D. degree in computer science from the University of California, Los Angeles in 2009. He is currently a Postdoctoral Research Fellow at the University of Bologna in Italy. His research interests include ad hoc, vehicular, and biomedical networks. More information can be found at <http://www.cs.unibo.it/~marfia>.



David Mack has studied Computer Science at Cambridge University and Mathematics and The Foundations of Computer Science at Oxford University. He is currently working in industry and exploring ideas for future doctorate study. More information about his profile and his work can be found at <http://www.edmack.com/>.



Dr. Giovanni Pau is a Research Scientist with Computer Science Department at the University of California, Los Angeles (UCLA). He received his Laurea Doctorate in computer science and the Ph.D. degree in computer engineering from the University of Bologna, Italy. His research interests include wireless ad hoc networks, Urban Sensing, and Distributed computing. He is the architect of the UCLA Campus Vehicular Testbed, an experimental facility to support urban sensing for pollution monitoring and Intelligent Transportation Systems. Since 2010

is leading the research at MPI-UCLA joint center on Ubiquitous Computing. More information can be found at <http://www.cs.ucla.edu/~gpau/>.



Dr. Cecilia Mascolo is a Reader in Mobile Systems in the Computer Laboratory, University of Cambridge. She has published extensively in the areas of mobile computing, mobility and social network modelling. She is and has been PI of various national and international projects on social network analysis, sensing and mobility modelling. Dr Mascolo has served as a programme committee member of many mobile, sensor systems and networking conferences and workshops. She has taught various courses on mobile and sensor systems and a Masters course on

social and technological network analysis. More details of her profile are available at www.cl.cam.ac.uk/users/cm542.



Prof. Mario Gerla received a graduate degree in engineering from the Politecnico di Milano in 1966, and his M.S. and Ph.D. degrees in engineering from UCLA in 1970 and 1973. He joined the Faculty of the Computer Science Department at UCLA, where he is now a professor. He worked for Network Analysis Corporation, New York, from 1973 to 1976. His research interests include distributed computer communication systems and wireless networks. He has designed and implemented various network protocols (channel access, clustering, routing, and transport) under DARPA and NSF grants. Currently, he is leading the ONR MINUTEMAN project at UCLA, with a focus on robust scalable network architectures for unmanned intelligent agents in defense and homeland security scenarios. He also conducts research on scalable TCP transport for the next-generation Internet (for recent publications, see <http://www.cs.ucla.edu/NRL>).