This is a preprint version of the following article, published by Elsevier (doi:10.1016/j.trc.2009.12.003):
T. Delot, N. Cenerario and S. Ilarri, "Vehicular Event Sharing with a Mobile Peer-to-Peer Architecture", Transportation Research Part C: Emerging Technologies, ISSN 0968-090X, 18(4):584-598, Elsevier, August 2010.

Vehicular Event Sharing With a Mobile Peer-to-Peer Architecture

Thierry Delot [a] , Nicolas Cenerario [a] , Sergio Ilarri [b]

[a] *LAMIH Laboratory, University of Valenciennes, Le Mont Houy, 59313 Valenciennes, France*

[b] *IIS Department, University of Zaragoza, Maria de Luna 1, 50018 Zaragoza, Spain*

Preprint submitted to Elsevier *3 December 2009*

Abstract

In vehicular ad hoc networks (VANETs), different types of information can be useful to drivers. Such networks are highly dynamic due to both the movements of the vehicles and the short range of the wireless communications. Thus, the information exchanges between vehicles about relevant information can only rely on short interactions. Therefore, an efficient mechanism to manage and disseminate the relevant information is required. Specifically, we present in this paper a system for data sharing in vehicular networks, which we call VESPA (Vehicular Event Sharing with a mobile Peer-to-peer Architecture). In this system, a new technique based on the concept of *Encounter Probability* is proposed for vehicles to share information using vehicle-to-vehicle communications. The objective is to facilitate the dissemination of information between vehicles when they meet each other, taking into account the relevance of the data to the drivers. Besides, the relevance must also be considered to inform a driver about the interesting events. Moreover, our proposal takes into account any type of event (e.g., available parking spaces, obstacles in the road, information relative to the coordination of vehicles in emergency situations, etc.) in the network. An experimental evaluation and the implemented prototype show the interest of the system.

Key words:
data management, event sharing, inter-vehicle communications, vehicular ad hoc networks

---

* Corresponding author. Tel.: +34 976 76 23 40. Fax: +34 976 76 19 14.
Email addresses: `Thierry.Delot@univ-valenciennes.fr` (Thierry Delot), `Nicolas.Cenerario@univ-valenciennes.fr` (Nicolas Cenerario), `silarri@unizar.es` (Sergio Ilarri).

1 Introduction

Nowadays, the car is indisputably the most heavily used mode of transportation. Unfortunately, its popularity has been accompanied by numerous problems, for example, in the areas of safety and the environment. To reduce the number of accidents, a variety of programs, generally involving “Intelligent Transport Systems”, have been initiated in Japan, Europe and the United States, attracting the interest of researchers both in academia and in industry. Thanks to the resulting research, Advanced Driver Assistance Systems (ADAS) were born. Some ADAS are already available on the market (navigation systems, warning systems to alert the driver when he/she is about to fall asleep in order to prevent him/her from crossing the center line, etc.), and many others are under development.

At the same time, wireless networks have proliferated, which has made inter-vehicle communications (IVC) possible. The primary objective of IVC is to improve road safety by developing new aids for drivers by exchanging information between vehicles in close proximity. For example, IVC could be used to inform drivers that an accident has occurred or that an obstacle has appeared on the road a few hundred meters ahead. As concerns data management for IVC, there is a fundamental difference with the driver assistance systems already available on the market, such as navigation systems. Thus, navigation systems only exploit static data, such as points of interest (for users to easily locate an airport, a restaurant, etc.), stored on a memory card. Although technologies such as Traffic Message Channel (TMC) (Klingenberg, 1993) can be integrated directly into a navigation system to deliver non-static data such as traffic and travel information (e.g., accidents) to drivers, IVC aims at sharing much more dynamic data between vehicles, both temporary and mobile events (emergency brakings, available parking spaces, etc.).

The wireless networks used for vehicles to communicate are short-range networks (about a hundred meters), which rely on standards such as IEEE 802.11 or Ultra Wide Band (UWB) (Luo and Hubaux, 2005) and provide bandwidth in the range of Mbps. Using such communication networks, a car can receive/send information (for example, about accidents, traffic congestion or available parking spaces) from/to its neighbors. Though in principle it would be possible to use wide-range networks, like the ones used for mobile telephony (GPRS/UMTS/...), they suffer from a lack of dynamism and a limited bandwidth (a few hundred Kbps). So, with such communication mechanisms, information about an event such as an emergency braking situation involving a preceding vehicle cannot be shared sufficiently early to allow a following vehicle to avoid a collision. Nevertheless, they could be used to provide access to distant services, such as web services delivering information about fuel prices.

Clearly, a dynamic inter-vehicle context with a very intermittent network connectivity (e.g., two vehicles within range of each other can move at high speeds in opposite directions) creates truly interesting data management challenges, which must be addressed to propose efficient driver assistance systems. In such dynamic environments, data access depends on inter-vehicle information exchanges. Data can be received from other vehicles and stored locally in a data cache. Then, query evaluation techniques can be used to sift through the stored information to determine what is relevant for that time and location, and issue a warning or transmit information to the driver when necessary[1]. Here, it is not trivial to determine how to control which data is relevant to be shown to the driver and which data should be transmitted to other vehicles. Both the spatial and temporal aspects of the information must indeed be verified. For example, in the case of information about an available parking space, an interested car must determine: 1) if it is close enough to the reported parking space, and 2) if the parking space was liberated recently enough and so it is probably still available. In this paper we propose VESPA, a system to share information about vehicular events in inter-vehicle ad hoc networks. The contributions of this work are summarized below:

- *Event classification*: A classification of the different events that may occur on the roads, based on direction and mobility features, is presented. Even mobile events, such as the appearance of an emergency vehicle (e.g., a police cruiser, an ambulance, or a fire engine) requiring other vehicles on the road to give way, are taken into account.
- *EP-based relevance evaluation*: A technique is proposed to evaluate the relevance of the information exchanged between vehicles by defining four-dimension mobility and direction vectors, computed for both vehicles and events. These vectors are then used to determine the *Encounter Probability* (*EP*) between a vehicle and an event.
- *Dissemination protocol*: An Encounter Probability-based technique to disseminate information in inter-vehicular ad hoc networks is presented.
- *Experimental evaluation*: The solutions presented in this paper have been strongly evaluated on a simulator. The results show the interest of the En-

---

[1] This obviously implies that not all the messages received by a vehicle are presented to the driver.

counter Probability both to evaluate the relevance of events for drivers and to disseminate the events among vehicles. Moreover, some tests have been performed with a real prototype.

The rest of this paper is organized as follows. Section 2 presents some related works. Section 3 proposes a classification of events relevant in inter-vehicle networks and describes the representation of events in VESPA. Section 4 explains how the Encounter Probability is calculated. Section 5 presents the proposed dissemination technique. Section 6 shows some experimental evaluations. Finally, Section 7 offers some conclusions and ideas for future research.

2 Related Work

Inter-vehicle communications is a recent field of research, with a main focus on data sharing. Naive data dissemination protocols like classical flooding have significant disadvantages; for example, in (Heinzelman et al., 1999) the problems of implosion, overlap and resource blindness are described. Similarly, geocasting protocols, such as (Navas and Imielinski, 1997), whose goal is to transmit data to all the targets within an area, are not appropriate in the context of vehicular ad hoc networks because they do not consider the features of the transmitted events and the mobility of vehicles. Therefore, data dissemination protocols specific to vehicular networks are being proposed.

Several works have studied some aspects that must be considered when developing a dissemination protocol in this context. For example, in (Adler and Strassberger, 2006) it is claimed that neither the size and shape of the dissemination area nor the message lifetime (amount of time that the message is kept alive in the network) should be set by the vehicle generating the event; instead, it should be adaptively determined (e.g., based on the current traffic pattern) by using a distributed approach. Besides, the importance of considering the relevance of events, called the *expected benefit* in (Adler et al., 2006; Eichler et al., 2006), especially when the bandwidth is scarce, is also emphasized. In the rest of this section, we describe some proposals in this area.

In the context of the *TrafficView* project (Nadeem et al., 2004, 2006), vehicles moving on roads with multiple lanes on each direction are considered. Specifically, in (Nadeem et al., 2006) three dissemination protocols for highway scenarios are compared analytically and experimentally: 1) dissemination by vehicles circulating in the same direction (*same-dir*), 2) dissemination by vehicles moving in the opposite direction (*opp-dir*), and 3) dissemination by vehicles moving in both directions (*bi-dir*). The analytical study considers two important elements: the *dissemination latency* (time needed to propagate data between two vehicles separated a certain distance) and the *broadcast utilization* (percentage of new area covered by a broadcast). In the experimental evaluation, the metrics evaluated are: the *knowledge percentage* (the percentage of vehicles that receive the data), the *accuracy* (in terms of the average error in estimating the location of other vehicles along the road, which is required in TrafficView), the *latency*, and the *utilization rate* (proportion of useful information received by the vehicles, as data about vehicles located behind on the road or outdated data are not useful). Based on the results of these experiments, it is claimed that the *opp-dir* model is the most efficient in general, except when the traffic in the opposite direction is sparse, in which case the *bi-dir* model is better[2]. The main limitations of this work are that only high-

---

[2] Other works, such as (Füβler et al., 2002), have also emphasized the importance

way scenarios are considered and that the relevance of data is not taken into account. Besides, different types of events are not explicitly considered.

In (Lochert et al., 2007) the focus is on urban environments. The authors explain that the network connectivity is a limiting factor for information dissemination, since chains of vehicles are needed to broadcast and a low traffic density may become a problem. So, the authors make a clear distinction between data transportation via locomotion[3] and via wireless communications, emphasizing the problem of lack of network connectivity that may occur depending on the density of equipped vehicles in an area. The main contributions of this work are the idea that *Stationary Supporting Units* (*SSUs*) are needed to alleviate this problem and different heuristics to decide the best locations to place them. By contrast, the data dissemination protocol of VESPA only needs vehicle-to-vehicle communications, although SSUs could be used to increase its coverage in areas with sparse traffic.

The *Mobi-Dik* project (Xu et al., 2004) proposes an *opportunistic exchange* mechanism (vehicles communicate with one another when they are close) inspired by the field of epidemiology. A vehicle with a certain piece of information acts as a disease carrier, and “contaminates” the nearby vehicles along its route. Once contaminated, these vehicles proceed to contaminate others. This dissemination principle is accompanied by mechanisms that monitor the relevance of the information (based on temporal and spatial criteria) in order to decide if it should be stored in the cache and/or broadcasted later on.

---

of the direction of vehicles for the dissemination protocol.

[3] That is, vehicles carry data to areas where they can be disseminated, which is usually called *store and forward* (Wu et al., 2004; Costa et al., 2006) or *carry and forward* (Zhao and Cao, 2006).

Although this mechanism is well adapted for cars to share information about available parking spaces (which is the case study for Mobi-Dik), it has not been designed to deal with other types of events (e.g., to relay information about an accident or an emergency braking situation).

In (Nittel et al., 2004), how to disseminate relevant geospatial information (road hazards) is studied. As in Mobi-Dik, an opportunistic exchange is proposed. According to this work, an information dissemination strategy must consider: the time needed since new data are available until the network stabilizes, the optimal distribution area around the event originator, and the time interval during which the data persist. Three different communication strategies are described in order to decide when a peer should communicate information to another peer in its neighborhood: the classical *flooding* strategy, the *epidemic* strategy (inform only a certain number of peers), and the *proximity* strategy (inform only the peers within a certain distance of the location of the event). These strategies are compared using different metrics: the *ignorance* (not knowing about a hazard found by a peer), the *redundancy* (receiving irrelevant information), and the number of messages transmitted. According to the experimental evaluation presented, the proximity strategy seems to achieve a good tradeoff. However, the need of more experiments is emphasized. Moreover, although the proposed dissemination strategies may be suited to specific events (mainly, road hazards), the features of different types of events (e.g., mobile events) are not explicitly considered.

The work in (Costa et al., 2006) proposes the use of a *propagation function* to decide the route that a message has to follow to reach a target spatial area. The originator of a message defines an appropriate propagation function (e.g., by considering the traffic conditions for the current time frame), which can be interpreted as a "gravitational field" where the message is attracted towards areas of minimum potential. The route traversed by the message is thus the result of evaluating the propagation function at each routing hop. On the basis of this propagation function, different dissemination approaches (both deterministic and probabilistic) are proposed and compared. The metrics used for performance evaluation are the *message delivery rate* (ratio between the nodes receiving the message and the total number of nodes, which should be closed to 100% within the target area and close to 0% otherwise) and the *network traffic* (total number of messages transmitted). However, how to define appropriate data propagation functions for different scenarios, which is a key element for the dissemination strategy, is not studied in the paper. Moreover, although the experimental results are promising, for both dense and sparse networks, more experiments with more mobility models (other than the Manhattan mobility model) are needed to represent realistic traffic patterns, as the authors themselves claim in relation to the proposed protocol *DFD-FSFD* (*Direction-aware Function Driven Feedback-augmented Store & Forward Diffusion*). An advantage of the approaches proposed is that, as in the dissemination protocol proposed in VESPA, there is no need to keep information about neighbors (which would be very expensive). It is also worth mentioning that the direction of a mobile node is considered in the *DFD-FSFD* protocol; this is also a factor used in VESPA to determine the relevance of events.

Finally, it is interesting to mention the work presented in (Kosch et al., 2006), which aims at achieving a scalable dissemination based on the relevance of the information exchanged. The relevance of a packet is computed based on the vehicle context (e.g., the current situation of the vehicular network), the message context (e.g., the amount of time elapsed since the message was generated), and the information context (e.g., the interest that other vehicles may have in the information). This work proposes to differentiate messages at two levels: in-vehicle message selection and inter-vehicle message selection. The goal of in-vehicle selection is to ensure that the messages are broadcasted ordered by relevance. Inter-vehicle message selection concerns modifications to Medium Access Control (MAC) schemas according to the relevance of the packets. Although the paper does not detail how the values of the different relevance parameters are set, it indicates that "the resulting parameters finally have to be weighted with application-dependent factors". Thus, the relevance function proposed is so general that the one proposed in VESPA could be seen as a specific instantiation based on the concept of Encounter Probability, whose usefulness will be emphasized along our work.

In relation to these works, VESPA provides several advantages. Thus, it considers different categories of events, including mobile events. Besides, it is a general approach (not focused only on highways or urban areas). Finally, it considers the relevance of events to disseminate data and also to inform the driver, based on the novel concept of Encounter Probability. The existing protocols and dissemination techniques mentioned above are interesting, but they mainly focus on how and when the information relative to events (accidents, obstacles, etc.) or resources (available parking slots, etc.) should be disseminated to other vehicles. Many existing solutions have been developed with particular types of events in mind, presented as an application example, and thus probably need different adaptations to manage other types of events efficiently. For instance, Mobi-Dik provides a very interesting solution to the problem of information-sharing inside a restricted spatio-temporal area. However, although the techniques proposed in Mobi-Dik are very well adapted for sharing information between cars about available parking spaces, they cannot be exploited to relay information about an accident or an emergency braking situation. On the contrary, VESPA is a generic system that can be deployed in cars to share and manage different types of events.

3 Events and Data Representation

Numerous types of events are possible in the context of inter-vehicle communications, since there is a lot of information that drivers may find relevant, about: accidents, traffic congestions, emergency braking situations, fuel prices, available parking spaces, emergency vehicles such as ambulances, obstacles in the road, or the behavior of drivers (e.g., strange maneuvers due to intoxication or lack of vigilance[4]), etc. Data about events occurring on the road or available resources (such as parking spaces) may have to be communicated to a potentially large set of vehicles, depending on the relevance of the data to the drivers. It should be noted that the generation of many events could be initiated using the numerous sensors embedded in modern cars (for example, by coupling the airbag system with the creation of an event representing an accident) or via other static data sources (e.g., sensors on a road)[5]. In the rest of this section, a system of event classification is proposed, and how these events are represented in VESPA is described.

---

[4] Lack of vigilance, or hypovigilance, can be detected today with oculometers using techniques that essentially count the number of eye blinks.

[5] This will prevent a driver from disseminating false information to his/her own benefit. Ensuring the reliability of messages manually generated by drivers is out of the scope of this paper and considered in works such as (Ostermaier, 2005).

3.1 Event Classification

We classify inter-vehicle network events in four categories: *stationary non-direction-dependent events*, *stationary direction-dependent events*, *mobile non-direction-dependent events*, and *mobile direction-dependent events.* In this classification, *direction-dependent events* are events that are not relevant for all nearby vehicles, but only for the vehicles traveling in a particular direction towards the event. On the other hand, *mobile events* are events whose locations change along time. Let us consider some examples. Available parking spaces correspond to stationary non-direction-dependent events, since they are static and may interest all vehicles close to that resource independently of their direction of movement. A warning about an accident is a stationary direction-dependent event because its location is fixed and only those vehicles that are expected to encounter the accident will find the message relevant, not the vehicles close to the accident but moving in the opposite direction. Messages warning vehicles of the lack of vigilance of a person driving on a two-way road are mobile non-direction-dependent events because they concern all vehicles likely to meet such driver, regardless of their direction of movement. Finally, an emergency vehicle broadcasting a message for other vehicles to yield the right of way is a mobile direction-dependent event.

It should be noted that, in our proposed classification, we group events according to mobility and direction features because these are the most relevant for the computation of the Encounter Probability (explained in Section 4).

In the following, we describe how the events are represented when they are created to be exchanged between vehicles. We distinguish several fields:

- *Key*. It is a unique value generated by combining the current time and the GPS location of the event with a randomly-generated sequence.
- *Version*. It is a number used to distinguish between different updates of the same event. Once generated, an event is disseminated among a set of potentially interested vehicles. To update the information transmitted (e.g., because a mobile event has moved), the vehicle which created the event may produce a new version of the same event.
- *Importance*. It is a value that helps to determine the urgency of presenting the information to the driver. An event with a high value for this field (e.g., an accident or an emergency braking) is expected to be relevant for any driver that may encounter the event, and so the driver must be informed. On the contrary, events with a low importance (e.g., available parking spaces) are reported to the driver only if he/she has requested such information.
- *CurrentPosition*. It is the location where the data was generated. Positions in VESPA correspond to GPS statements and thus include three-dimensional coordinates as well the GPS time. Using the GPS time allows to avoid synchronisation problems between the clocks of the different vehicles.
- *DirectionRefPosition* and *MobilityRefPosition*. They are two preceding positions that provide each vehicle receiving the event with information needed to compute the mobility and direction of the event, which is necessary to estimate its relevance (see Section 4). It should be noted that, as an alterna-

tive, a vehicle could just send periodically the location of the event and the receiving vehicle could estimate its direction based on the locations received. Explicitly communicating a representative reference position simplifies this task and reduces the network overload, similarly to other tracking proposals in the field of moving object databases (e.g., see (Wolfson et al., 1999)).

- *LastDiffuserPosition.* It is the position of the last vehicle which relayed the message. This value is used by the dissemination protocol (see Section 5).
- *HopNumber.* It is the number of rediffusions of the message (see Section 5).
- *Description.* It describes more precisely the represented event, so it allows to transmit concrete information to drivers when they need to be warned.

It should be noted that the category of the event (stationary or mobile, direction-dependent or not) is not explicitly represented as an attribute of the event, as it can be inferred from other message fields.

4 Computing the Relevance of Events

One of the major problems in inter-vehicle communications is determining the relevance of an event to a specific receiving vehicle. In this section, it is first shown how the data describing an event is exploited to compute four-dimensional mobility and direction vectors (three dimensions for the spatial coordinates and one for the temporal dimension). Then, it is introduced the notion of *Encounter Probability*, which is used to estimate whether a vehicle is expected to encounter an event in order to decide if the event needs to be considered (Delot et al., 2008).

4.1 Mobility and Direction Vectors

To compute the relevance of an event for a specific vehicle, it is necessary to have an estimation of the trajectory of the vehicle (and, in the case of a mobile event, also an estimation of the trajectory of the event) in order to obtain the probability that the vehicle meets the event. The direction of a moving object (a mobile event or a vehicle) is modeled by using vectors that run between a preceding reference position and the current position of the object. These vectors are computed thanks to GPS position statements obtained regularly. It should be noted that although VESPA relies on a positioning system to obtain the positions needed to generate the vectors, neither a navigation system nor any kind of roadmap is used. On the contrary, works such as (Adler and Strassberger, 2006; Mammar et al., 2005; Leontiadis et al., 2009) propose the use of road maps. Such maps would allow future vehicle positions to be predicted more easily. So, if digital maps were available, the accuracy could improve because they could be used to make better predictions. However, our proposal is general enough to work without such maps.

An estimated future position depends on the time interval selected between the position statements used to compute the vectors. Thus, if the interval is large, the estimation of the future position is not precise but provides an overall impression of the direction of the object. If the time interval is shorter, then the estimation is much more precise on the short term but no global view of the movement can be observed. The role of the *mobility vector* is to provide an overall impression of the movement of the object in addition to a good estimated future position, so an “average” interval must be used to compute the mobility vector (see arrow C in Figure 1). On the contrary, the *direction* *vector* (arrow B in Figure 1) is computed with a short interval, and therefore provides a quite precise estimated future position but only in the short term.

[place Fig. 1 about here]

For simplicity, to compute these vectors we use cartesian coordinates, which are obtained from the GPS coordinates through a planar projection[6]. Besides its own mobility and direction vectors, each vehicle can compute the *direction and mobility vectors of the events* it receives. For that purpose, it uses the data associated to events, and more precisely the *CurrentPosition* attribute and either the *DirectionRefPosition* or the *MobilityRefPosition* attribute, respectively.

For each event, the *mobility and direction vectors of the vehicle in relation to the event* are computed by changing the frame of reference. Figure 2 illustrates this change. The vectors of one vehicle and one event are represented on the left side of the figure, and the resulting vector after the frame of reference has been changed is shown on the right side. This simplifies the computation of the Encounter Probability (Section 4.2) by allowing a single vector for each couple <vehicle, event> to be managed, regardless of the category of the event.

[place Fig. 2 about here]

To explain how this change of frame of reference is computed, the mobility vector of a vehicle A between t_{A_1} and t_{A_2} (Δt_A) and that of an event B between t_{B_1} and t_{B_2} (Δt_B) is considered:

---

[6] The most suitable type of projection depends on the region of the earth.

$$V_A(t_{A_1}, t_{A_2}) = \begin{pmatrix} x_A \\ \\ y_A \\ \\ z_A \\ \\ \Delta t_A \end{pmatrix}, V_B(t_{B_1}, t_{B_2}) = \begin{pmatrix} x_B \\ \\ y_B \\ \\ z_B \\ \\ \Delta t_B \end{pmatrix}$$

The first step is to modify the mobility vectors in order to manage the same time basis (fourth dimension) for both vectors:

$$V'_A(t_{A_1}, t_{A_2}) = V_A(t_{A_1}, t_{A_2}) \times \Delta t_B$$

$$V'_B(t_{B_1}, t_{B_2}) = V_B(t_{B_1}, t_{B_2}) \times \Delta t_A$$

Then, the mobility vector of vehicle A in relation to event B is computed by subtracting $V'_A(t_{A_1}, t_{A_2})$ and $V'_B(t_{B_1}, t_{B_2})$:

$$V_{AB}(t_{A_1}, t_{A_2}, t_{B_1}, t_{B_2}) = \begin{pmatrix} (x_A \times \Delta t_B) - (x_B \times \Delta t_A) \\ \\ (y_A \times \Delta t_B) - (y_B \times \Delta t_A) \\ \\ (z_A \times \Delta t_B) - (z_B \times \Delta t_A) \\ \\ \Delta t_A \times \Delta t_B \end{pmatrix}$$

In the case of a stationary event, for which the mobility vector of the event is the *null* vector, $V_{AB}(t_{A_1}, t_{A_2}, t_{B_1}, t_{B_2})$ is equal to $V_A(t_{A_1}, t_{A_2})$.

The mobility and direction vectors of a vehicle in relation to an event can be used to compute an Encounter Probability to estimate whether the vehicle will meet the event, as it is described in the following section.

4.2 Encounter Probability

In order to estimate if an event is relevant or not, we use the concept of *Encounter Probability.* In this section, we first introduce this concept, and then we describe how the different parameters involved in its computation can be set.

4.2.1 Computation and Use of the Encounter Probability

The Encounter Probability is based on the following four elements:

- The minimal geographical distance between the vehicle and the event over time (Δd).
- The difference between the current time and the time when the vehicle will be closest to the event (Δt).
- The difference between the time when the event is generated and the moment when the vehicle will be closest to the event (Δg).
- The angle between the direction vector of the vehicle and the direction vector of the event (represented by a colinearity coefficient c).

As an example, the geometrical representation of Δd and Δt in a certain scenario is shown in Figure 3. In this example, a stationary event is considered (to facilitate the graphical representation of the mobility vector in relation to the event), but the principle would be the same with a mobile one. B represents the position of the vehicle, C is the position of the event, and $\overrightarrow{AB}$ is the mobility vector of the vehicle in relation to the event. The closest point to C on the straight line between A and B (point D) can then be determined, which allows a right-angled triangle to be constructed in D with $[BC]$ as hypotenuse.

$|DC|$ ($= \Delta d$) represents the minimal geographical distance between the vehicle and the event over time, and $|BD|$ is the distance between the vehicle and the point D. Since the mobility vector $\overrightarrow{AB}$ has a temporal dimension, $|BD|$ can be converted into time to obtain Δt.

[place Fig. 3 about here]

The *age of the event* (Δg) corresponds to the interval between the moment when the event was generated and the moment when the vehicle is closest to the event. It is obtained from the message received, by using the event generation time (stored in *CurrentPosition*), and the Δt previously computed.

Finally, as mentioned previously, the vehicle is able to estimate its direction vector and the direction vector of the event. These two vectors produce a *colinearity coefficient* (c), which is a measure of the angle formed by the vectors. For direction-dependent events, it allows to determine if the directions of the vehicle and the event match. For non-direction-dependent events, c is 0.

The *Encounter Probability* (0%-100%) estimates the probability that the vehicle meets an event as follows:

$$EP = \frac{100}{\alpha \times \Delta d + \beta \times \Delta t + \gamma \times \Delta g + \zeta \times c + 1}$$

where α, β, γ and ζ are *penalty coefficients* with non-negative values which are used to balance the relative importance of the Δd, Δt, Δg, and c values. The bigger the coefficient is, the more penalized the associated valued is when computing the Encounter Probability. Thus, the greater the α value, the shorter the spatial range where the event is relevant; β and γ are used to consider only the most recent information and the information about events that will be encountered quickly; finally, ζ is used to weigh the importance of the colinearity of the direction vectors for direction-dependent events.

The Encounter Probability of an event is compared with two different thresholds: the relevance threshold and the storage threshold. The *relevance threshold* indicates the minimum value of the EP to assume that the vehicle is likely to meet the event, which means that the event is considered *relevant for the vehicle*. The *storage threshold* (no higher than the relevance threshold) indicates the minimum value of the EP to consider an event received as *potentially relevant for the vehicle*. A potentially relevant event is stored on the vehicle so that the vehicle can watch it closely. Thus, such an event may be considered not relevant enough by the vehicle at that moment but it could become relevant in the near future (e.g., because the vehicle modifies its direction). In this way, the vehicle can also simply act as a carrier of the event. When a vehicle receives an event, the event is not considered if its EP is lower than the storage threshold (the potential relevance of the event for the vehicle is not enough). Otherwise, the vehicle stores the event and, additionally, a warning will be communicated to the driver if and only if: 1) the EP is also greater than the relevance threshold, and 2) either the *Importance* of the event is high (e.g., it is an accident) or the driver has specified his/her interest in that type of event (i.e., *the event is relevant for the driver*).

Thus, a distinction between events *relevant for the vehicle* and events *relevant for the driver* is made. An event is relevant for the vehicle if the vehicle will probably meet the event. An event is relevant for the driver if it is relevant for his/her vehicle and besides the driver is interested in that type of event. Decisions about storing and disseminating events are based on the relevance for the vehicle, without considering the interests of the driver. In this way, VESPA emphasizes the cooperation among vehicles (which could be encouraged by using mechanisms such as those presented in (Wolfson et al., 2004)).

4.2.2 Choice of the Penalty Coefficients

The α, β, γ and ζ penalty coefficients are used to balance the relative importance of the $\Delta d, \Delta t, \Delta g$ and c components. Considered individually, they allow to define bounds on the relevance of events. For example, if the relevance threshold is set to 75% for the Encounter Probability, a value of $\alpha \geq \frac{1}{300}$ implies that if the minimum geographical distance between the vehicle and the event over time (Δd) is bigger than 100 meters, then the event will be considered as not relevant independently of the values of the other parameters:

$$75 \leq \frac{100}{(\alpha \times 100+1)} \Rightarrow \alpha \leq \frac{1}{300}$$

In the same way, β sets a maximum time interval between the current position of the vehicle and the position of the vehicle when it is closest to the event; if this interval is exceeded, the event is considered not relevant. For example, for values of $\beta \geq \frac{1}{900}$ an event will not be considered relevant if the time elapsed when the vehicle is at the closest distance from the event is five minutes or more. Similarly, γ is used to penalize the relevance according to the age of the event. In practice, γ should be set according to the frequency used to generate new versions of potentially long-term events (e.g., if this period is five minutes, then it is possible to set $\gamma = \frac{1}{900}$). Finally, ζ may induce a maximum tolerance on the angle between the direction vectors for direction-dependent events.

The importance of $\Delta d, \Delta t, \Delta g$ and c depends on the event considered. For instance, a message describing a traffic congestion should be broadcasted several kilometers away from the place where it is located, for drivers to have the opportunity to change their itinerary; the penalty on Δt should so be very low. On the other hand, when dealing with parking spaces, the penalty on Δt should be more important because a driver is only interested in nearby parking spaces. Moreover, for the same type of event, the penalty coefficients could be modified according to the current time and date. For example, when dealing with parking spaces in urban areas, the penalty on the age Δg should probably be more penalizing on Saturday afternoons than on Monday nights.

5 Dissemination Protocol

So far, the concept of Encounter Probability has been considered to determine whether the information contained in a message received on a vehicle should be stored and/or a warning for the driver should be produced. In this section, it will be shown that the EP is also useful to disseminate the information about events between vehicles (Cenerario et al., 2008). First, we will present a technique by which the dissemination area is adapted dynamically to ensure that each vehicle for which certain information is interesting will receive that information. Then, we will explain how the network usage is minimized.

5.1 Adaptive Dissemination Area

The proposed protocol should be able to disseminate different types of events in the inter-vehicle network. Therefore, different dissemination modes have to be supported. For example, the information about an accident has to be diffused only to the vehicles driving in its direction. Similarly, information about an emergency braking, like an accident, has to be diffused to the vehicles driving in a particular direction. Nevertheless, whereas the information should be relayed far away from the place where the accident took place in the first case, it should only be relayed a few hundred meters away in the case of an emergency braking. As a final example, information about an available parking space has to be transmitted to all close vehicles (the direction does not matter), as it may interest them.

The use of the EP to determine the vehicles which have to rediffuse certain information they received allows to disseminate the messages in the right direction, that is, towards the vehicles for which these messages may be relevant. This also ensures that the information about an event is maintained, during the dissemination phase, close enough to be relevant. For instance, the information about an available parking space would not be interesting for persons driving several kilometers away from it. Therefore, in the proposed dissemination solution, each time a vehicle receives a message, it computes the EP for the corresponding event. If the value obtained is bigger than a certain *diffusion threshold*, it has to rediffuse the message. Otherwise, it does not consider the message. Thus, while the event is considered relevant by a vehicle in a particular area, it is relayed to the neighboring vehicles, and so on.

Moreover, the EP also avoids the dissemination of obsolete events. The information diffused for events such as a parking space or an emergency braking is only relevant for a short period of time. For events with a longer lifetime (e.g., an accident), it is possible to adapt the value of the corresponding penalty coefficient γ in order not to penalize too much the EP with the age of the event. Anyway, for long-duration events, new versions of the same event have to be produced to continue informing the arriving vehicles and update the information about the event (including its mobility profile). For example, in case the car of a sleepy driver is blocked due to a traffic congestion, new versions of the generated mobile non-direction-dependent event (warning about his/her lack of vigilance) would indicate the speed reduction. If the car should stop for a while, for the same reason, then a new version of the event would include two reference positions equal to the current position and then represent a temporary stationary event.

Motivated by the use of the Encounter Probability, revocation/invalidation messages indicating that an event has disappeared are not considered in VESPA. Instead, each vehicle estimates when the event becomes irrelevant. Considering explicit revocation messages would be challenging because it is very difficult to guarantee that all the vehicles previously informed of an event could be reached by an event revocation message. When revocation messages are not used, the number of messages exchanged in the vehicular network is also minimized. Thus, in VESPA revocation messages are not generated, although it could be useful to release such messages in certain circumstances when the conditions have changed significatively (e.g., an important accident that has been cleared up). It should also be emphasized that not all the events/messages received by a vehicle are shown to the driver (a query evaluation is performed on the data stored locally, as mentioned in Section 1). Therefore, when an event disappears the vehicles that received that event previously may have never presented such information to their driver.

5.2 Limited Bandwidth Use

To avoid flooding, and so network congestions, the proposed solution aims at desynchronizing the rediffusions, similarly to other approaches such as *contention-based forwarding* (Füβler et al., 2003), proposed for unicast message delivery in vehicular networks. Since the value of the EP may be greater than the diffusion threshold used for many vehicles, it is necessary to limit the number of diffusions of a single message. Therefore, each vehicle will wait for a period T before rediffusing the message. If the vehicle receives the same message before that period elapses, then it must not rediffuse it again because another vehicle already performed the rediffusion. The duration of that period T depends on the distance between the receiving vehicle and the one which sent the message. The intuition behind this is to choose, among the neighbors which received the message, the farthest neighbor from the sender to relay the message. Indeed, this farthest neighbor may have the greatest number of neighboring vehicles not yet informed about the event being transmitted. It is so the best candidate to try to broadcast the message to all concerned vehicles as quickly as possible. The value of T is determined by each vehicle as follows:

$$T = D \times (1 - \tfrac{d}{r})$$

where D is the maximum time to wait before rediffusing, r is the communication range of the wireless network used by the vehicles (e.g., 200-400 meters), and d corresponds to the distance between the receiving vehicle and the diffusing vehicle. The value of d is computed using the *LastDiffuserPosition* attribute in the message. Since d may vary from 0 to r, T is between 0 and D.

This approach allows a message to propagate far from the origin (if needed) and, at the same time, minimizes the number of duplicated messages received by the vehicles. An alternative where the sending vehicle decides which neighbor should rediffuse would be unrealistic (e.g., see (Füβler et al., 2003)), as this would require that the vehicles track the locations of their neighbors, which will be changing continuously.

Upon reception of a message, a vehicle has to avoid rediffusing such a message if it has rediffused it already before. It may happen also that none of the vehicles receiving a message about an event determines it relevant enough, after computation of the EP, to rediffuse it. Similarly, it may occur that there is no vehicle receiving the message because no vehicle is within range of the sending vehicle. To overcome these situations, the following approach is used. Each time a vehicle diffuses a message, it waits during D seconds. Then, if it did not receive the message during that interval, it will periodically resend it until another vehicle estimates the event relevant and so diffuses it also (or until the event becomes irrelevant for the vehicle). The goal is to keep the message *alive* while it is relevant. The *HopNumber* attribute (see Section 3) is increased every time a vehicle relays a message, and it is used to determine if a message received is a rediffusion (the *HopNumber* of the message corresponds to the *HopNumber* of an event with the same *Key* diffused previously by the receiving vehicle) and therefore should not be rediffused again.

6 Experimental Evaluation

In this section, some experimental results are presented. Even though we have developed a working prototype of VESPA[7], for obvious scalability reasons (e.g., need of several cars with appropriate mobile devices), the proposed relevance estimation and dissemination techniques could not be evaluated in a complex scenario with the real prototype. Therefore, experiments using a

---

[7] See `http://www.univ-valenciennes.fr/ROI/SID/tdelot/vespa/`

simulator were performed. In this section, the experimental settings are first described. Then, the usefulness of the Encounter Probability as a metric to estimate the relevance of the events is evaluated. Finally, the proposed dissemination protocol is tested.

6.1 Experimental Settings

The performance and feasibility of applications and protocols for vehicular networks are usually tested through simulation. Two important components are considered to simulate vehicular networks: the mobility of the vehicles and the communication network. Regarding the simulation of vehicles, we need to consider appropriate mobility models (Haerri et al., 2009). In this sense, there are several available traffic simulators, both commercial –such as *CORSIM* (`http://mctrans.ce.ufl.edu/featured/tsis/version5/corsim.htm`) or *TRANSIMS* (`http://transims.tsasa.lanl.gov/`)– and free –such as *VanetMobiSim* (Fiore et al., 2007) or *GrooveNet* (Mangharam et al., 2006)–. On the other hand, different network simulators have been developed, such as *NS2* (`http://www.isi.edu/nsnam/ns/`), *GloMoSim* (`http://pcl.cs.ucla.edu/projects/glomosim/`), or *JiST-SWANS* (`http://jist.ece.cornell.edu/`). Some works also propose an integrated solution, such as *TraNS* (Piórkowski et al., 2008) (`http://trans.epfl.ch/`).

Choosing the simulator/s to use is a big challenge because there is a wide variety of existing simulators. Moreover, configuring and adapting a generic simulator to one's needs is frequently an overwhelming task, especially if the goal is to test data sharing mechanisms rather than network protocols. The chosen simulator should allow to represent the environment, define different routes for vehicles, simulate vehicles with no specific destination (e.g., searching for a parking space), and integrate various speeds and traffic conditions. Besides, when an event occurs it is necessary to know which vehicles encountered that event and when. Moreover, the behavior of the drivers must be affected when certain events are received (e.g., an event indicating a traffic jam could encourage a driver to change his/her route). Finally, all the necessary information to evaluate the proposed dissemination and relevance mechanisms should be collected. Different alternatives were studied, but it proved difficult to implement and evaluate the proposed system using an existing simulation tool. So, a new simulator adapted to the evaluation needs of VESPA was developed.

The developed simulator allowed to simulate realistic contexts. To represent different kinds of curves, some segments of the roads are represented using Bezier curves. This is very important in the experiments presented in this paper, since one major goal is to validate the use of mobility vectors on different kinds of roads (with different curve profiles). Different mobility models are implemented in the simulator; specifically, in the experiments that are presented in this paper the vehicles follow the shortest routes towards random target locations. Messages are then transmitted between vehicles. When a message is sent, all the close-enough vehicles receive it (according to the considered communication range), and then they have to decide if they should store the event, inform the driver, and/or broadcast the message according to the value of the Encounter Probability. Experiments using the VESPA prototype in a real environment provided useful information to calibrate the simulator.

The results presented in the following correspond to simulations performed using two different road configurations representing segments of highways or main roads (see Figure 4). The first one gives a more precise view whereas the second one corresponds to a part of the route between the cities of Valenciennes and Lille (in France). Six stationary events, located in strategic places, were considered to evaluate the proposal. All the events are direction-dependent, except *Event5*, which is non-direction-dependent. These events do not disappear during the experiments. Between 500 and 600 vehicles, with random trajectories, are simulated. Their speed usually varies between 90 km/h and 110 km/h, depending on the road (e.g., vehicles slow down in a curve).

[place Fig. 4 about here]

Through experimentations, we chose values for the penalty coefficients that provide good results independently of the road configuration: $\alpha = 0.0033$, $\beta = 0.0010$, $\gamma = 10^{-8}$, and $\zeta = 0.25$. Naturally, it is possible to refine these coefficients according to the type of event or the specific context (as explained in Section 4.2.2). Similarly, based on experiments, it was decided to compute mobility vectors using position statements performed every 500 meters and direction vectors using position statements measured every 30 meters. The relevance threshold and the diffusion threshold are both set to 75%. The value of D for the dissemination protocol was set to one second and the communication range r considered is 200 m. The time needed to send a message from one vehicle to another within its communication range is set to 200 ms plus an extra delay between zero and one second depending on the distance between the vehicles; these are pessimistic values, which will allow us to evaluate the benefits of VESPA in a wireless environment with high communication delays.

The objective of the first set of simulations is to evaluate the usefulness of the Encounter Probability to estimate the relevance of events. It is assumed in these experiments that all the events are diffused using a flooding-based strategy. The use of such a flooding-based technique to disseminate messages would not be appropriate for deployment on real vehicles, but it allows to separate the evaluation of the usefulness of the Encounter Probability as a relevance measure from the evaluation of the proposed dissemination protocol.

In the simulations, the percentage of vehicles which estimated an event relevant before meeting it was measured. The vehicles that do not meet the event are not considered in the computation of this percentage. However, the experiment that we present in the following shows that all the vehicles meeting the event receive the corresponding message in advance. In Figure 5, the evolution of warned vehicles as a function of the reaction time available before meeting the event is presented (no vehicle travels on the road of *Event6*, and so it does not appear in the figure since, as desirable, it is never reported).

[place Fig. 5 about here]

For *Event1*, located on a straight road just after a sliproad, all the vehicles which encountered the event presented a warning to the driver and more than 93% had transmitted that warning 30 seconds before meeting the event. Although many vehicles generated a warning even if they never met the event, this was expected. Thus, since the event is very close to the intersection of two roads, the vehicles which turned on the right before encountering *Event1* had already generated a warning. Fortunately, those vehicles re-evaluated *Event1* as no more relevant very quickly after turning right (after about two seconds).

Event2 is located on a very strong curve. Despite the difficulty of this situation, the percentage of vehicles which transmitted a warning to the driver remains very high. The drivers reduced their speed in the curve, which gave them extra time to react (although the warnings were generated later than in other cases).

The results obtained for *Event3* confirm the good results obtained for *Event1*. However, the results obtained for *Event4* are not as good, with a warning generated only four seconds before meeting the event for most vehicles. This event is located after a strong change of direction, which is a very difficult situation. Thus, the system cannot predict with enough anticipation that the vehicle will meet the event after the strong curve. Therefore, the Encounter Probability computed by the vehicles reached the relevance threshold of 75% late. Fortunately, this situation can be improved by using *adaptive vectors*, which implies choosing the reference positions used to compute the mobility and direction vectors according to the distance separating the vehicle and the event. This allows to limit the impact of the direction changes by improving the accuracy of the estimated direction as the vehicle gets closer to the event. In this way, the warning is generated about eight seconds in advance (so the drivers have more time to react). This is shown in Figure 6, where adaptive vectors are also examined for *Event2* (the other event located on a strong curve). It can be concluded that adaptive vectors can be very useful in situations with strong changes of direction; however, they also imply a higher processing cost because they must be recomputed constantly.

Finally, *Event5* is considered relevant earlier than *Event2*, *Event3*, and *Event4*, as the direction does not penalize for non-direction-dependent events.

[place Fig. 6 about here]

To summarize, the results obtained are rather satisfactory, since the vehicles presented a warning to the driver in time, in every configuration.

6.3 Evaluation of the Dissemination Protocol

In the following, the results of the simulations performed to evaluate the proposed dissemination protocol based on the Encounter Probability are presented. Through experimentation, we first observed that all the vehicles received the relevant events before meeting them. In the experiments presented in this section, two metrics will be considered: the number of messages emitted, and the time needed to deliver a message to a vehicle. The proposed protocol will be compared with two alternatives: flooding and periodic flooding.

First, we measure the total number of messages diffused at each second since the generation of an event. Figure 7 is for low traffic density (about one vehicle every 100 meters) and Figure 8 for a a scenario with high traffic density (about one vehicle every 10 meters).

[place Fig. 7 about here]

[place Fig. 8 about here]

In the figures, we can first observe that the proposed dissemination solution strongly limits the number of messages exchanged. Regarding the evolution of messages transmitted it can be noticed that, with both the proposed dissemination mechanism and traditional flooding, the number of messages decreases after some seconds because the *tail* of the diffusion chain has been reached (the event is not considered relevant enough by the farthest vehicles to be diffused again). However, with the approach based on traditional flooding the number of messages reaches zero after only a few seconds, and so the information about the event stops propagating. Periodic flooding (where the messages are re-emitted periodically) could be an alternative, but it has a very high cost in terms of messages transmitted. With the proposed dissemination approach, the broadcast of the message continues at the *tail* of the diffusion chain (even if it cannot be seen clearly in Figure 8 due to the scale, a few messages are still diffused after ten seconds) as long as the EP is high enough. Therefore, the previous problem is solved by the proposed dissemination protocol.

Since the proposed solution may introduce waiting times at each hop of the dissemination process, another goal of the experiments was to evaluate this additional cost. Thus, Figure 9 (for low traffic density) and Figure 10 (for high traffic density) show, for the different dissemination approaches, the time needed for the vehicles to receive the information according to the distance separating those vehicles from the event. Using the proposed dissemination mechanism, the vehicles receive the information about the event slightly later. Nevertheless, the additional cost is limited, even in the worst case (i.e., when the traffic is low), and enough time remains available for the driver to react according to the information transmitted to him/her.

[place Fig. 9 about here]

[place Fig. 10 about here]

Summing up, according to our experiments, the proposed dissemination protocol keeps the number of message transmissions limited and at the same time is able to keep the information about an event alive in the network as long as necessary (i.e., until the event stops being relevant).

7 Conclusion and Perspectives

In this paper we have presented VESPA, a system for sharing data about different types of events in a vehicular ad hoc network (VANET). VESPA is thus complementary to existing navigation systems supporting only static information such as points of interest (POI), since it allows drivers to be informed of ephemeral events occurring on the roads. The proposal is based on the concept of *Encounter Probability* (*EP*), which is computed both to estimate the relevance of the events and to disseminate the events in the vehicular network. By considering the EP, the system is able to manage different types of events transparently. Moreover, the experimental results are really promising. On the one hand, the usefulness of the Encounter Probability to estimate the relevance of an event has been shown. On the other hand, our experiments also indicate that the drivers receive the interesting events well in advance while, at the same time, the cost of the dissemination protocol is limited.

Works such as (Young et al., 2003; Amditis et al., 2006; Bach et al., 2008), highlight the necessity to analyze carefully the impact of in-vehicle technology on drivers. We have performed some usability tests of our prototype, but more work is needed, for example to consider possible conflicts (Tuttlies et al., 2007) between different in-vehicle services. We have initiated a cooperation with the *SHM (Systemes Homme-Machine)* and PERCO (*Psychologie et Ergonomie Cognitive)* groups at the *Laboratoire d'Automatique, de Mecanique et d'Informatique Industrielles et Humaines* at the University of Valenciennes to analyze these aspects in more detail. As a different line of future work, we are also investigating ways to use the data exchanged between vehicles to generate knowledge that can be exploited later by the vehicles (Defude et al., 2008).

Acknowledgements

This work was partly supported by the Nord-Pas-de-Calais region, in the context of the D4S (Design for Safety) project, and the CICYT project TIN2007-68091-C02-02. We thank the anonymous reviewers for their insightful and useful comments. Last but not least, detailed and interesting comments by Raquel Trillo and Eduardo Mena helped us to refine this paper.

References

Adler, C., Eichler, S., Kosch, T., Schroth, C., Strassberger, M., September 2006. Self-organized and context-adaptive information diffusion in vehicular ad hoc networks. In: Third International Symposium on Wireless Communication Systems (ISWCS'06). IEEE Computer Society, pp. 307–311.

Adler, C., Strassberger, M., October 2006. Putting together the pieces – a comprehensive view on cooperative local danger warning. In: 13th ITS World Congress and Exhibition on Intelligent Transport Systems and Services (ITS'06).

Amditis, A., Polychronopoulos, A., Andreone, L., Bekiaris, E., September 2006. Communication and interaction strategies in automotive adaptive interfaces. Cognition, Technology and Work 8 (3), 193–199.

Bach, K. M., Jæger, M. G., Skov, M. B., Thomassen, N. G., 2008. You can touch, but you can't look: interacting with in-vehicle systems. In: 26th

SIGCHI Conference on Human Factors in Computing Systems (CHI'08). ACM, pp. 1139–1148.

Cenerario, N., Delot, T., Ilarri, S., June 2008. Dissemination of information in inter-vehicle ad hoc networks. In: Intelligent Vehicles Symposium (IV'08). IEEE Computer Society, pp. 763–768.

Costa, P., Frey, D., Migliavacca, M., Mottola, L., September 2006. Towards lightweight information dissemination in inter-vehicular networks. In: Third International Workshop on Vehicular Ad Hoc Networks (VANET'06). ACM, pp. 20–29.

Defude, B., Delot, T., Ilarri, S., Martini, J. Z., Cenerario, N., July 2008. Data aggregation in VANETs: the VESPA approach. In: First International Workshop on Computational Transportation Science (IWCTS'08), in conjunction with MobiQuitous'08. ICST.

Delot, T., Cenerario, N., Ilarri, S., April 2008. Estimating the relevance of information in inter-vehicle ad hoc networks. In: First International Workshop on Sensor Network Technologies for Information Explosion Era (SeNTIE'08), in conjunction with MDM'08. Workshop proceedings, pp. 151–158.

Eichler, S., Schroth, C., Kosch, T., Strassberger, M., July 2006. Strategies for context-adaptive message dissemination in vehicular ad hoc networks. In: Second International Workshop on Vehicle-to-Vehicle Communications (V2VCOM'06), in conjunction with MobiQuitous'06. IEEE Computer Society, pp. 1–9.

Fiore, M., Haerri, J., Filali, F., Bonnet, C., March 2007. Vehicular mobility simulation for VANETs. In: 40th Annual Simulation Symposium (ANSS'07). IEEE Computer Society, pp. 301–309.

Füβler, H., Mauve, M., Hartenstein, H., Ksemann, M., Vollmer, D., July 2002. A comparison of routing strategies for vehicular ad hoc networks. Tech. Rep.

TR-02-003, Department of Computer Science, University of Mannheim.

Füβler, H., Widmer, J., Kasemann, M., Mauve, M., Hartenstein, H., November 2003. Contention-based forwarding for mobile ad hoc networks. Ad Hoc Networks 1 (4), 351–369.

Haerri, J., Filali, F., Bonnet, C., 2009. Mobility models for vehicular ad hoc networks: a survey and taxonomy. IEEE Communications Surveys and Tutorials, to appear.

Heinzelman, W. R., Kulik, J., Balakrishnan, H., August 1999. Adaptive protocols for information dissemination in wireless sensor networks. In: Fifth ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'99). ACM, pp. 174–185.

Klingenberg, W., October 1993. Location referencing systems for dynamic route guidance applications. In: Vehicle Navigation and Information Systems Conference. IEEE Computer Society, pp. 441–444.

Kosch, T., Adler, C. J., Eichler, S., Schroth, C., Strassberger, M., October 2006. The scalability problem of vehicular ad hoc networks and how to solve it. IEEE Wireless Communications, 22–28.

Leontiadis, I., Costa, P., Mascolo, C., March 2009. Persistent content-based information dissemination in hybrid vehicular networks. In: Seventh IEEE International Conference on Pervasive Computing and Communications (PerCom'09). IEEE Computer Society.

Lochert, C., Scheuermann, B., Caliskan, M., Mauve, M., January 2007. The feasibility of information dissemination in vehicular ad-hoc networks. In: Forth Conference on Wireless On demand Network Systems and Services (WONS'07). IEEE Computer Society, pp. 92–99.

Luo, J., Hubaux, J.-P., 2005. A survey of research in inter-vehicle communications. In: Embedded Security in Cars – Securing Current and Future

Automotive IT Applications. Springer, pp. 111–122.

Mammar, S., Nouvelière, L., Netto, M., Glaser, S., November 2005. WILL-WARN hazard detection module and relevance check procedure. In: 12th ITS World Congress and Exhibition on Intelligent Transport Systems and Services (ITS'05).

Mangharam, R., Weller, D., Rajkumar, R., Mudalige, P., Bai, F., July 2006. GrooveNet: A hybrid simulator for vehicle-to-vehicle networks. In: Second International Workshop on Vehicle-to-Vehicle Communications (V2VCOM'06), in conjunction with MobiQuitous'06. IEEE Computer Society, pp. 1–8.

Nadeem, T., Dashtinezhad, S., Liao, C., Iftode, L., July 2004. TrafficView: Traffic data dissemination using car-to-car communication. ACM SIGMOBILE Mobile Computing and Communications Review, Special Issue on Mobile Data Management 8 (3), 6–19.

Nadeem, T., Shankar, P., Iftode, L., July 2006. A comparative study of data dissemination models for VANETs. In: Third International Conference on Mobile and Ubiquitous Systems (MobiQuitous'06). IEEE Computer Society, pp. 1–10.

Navas, J. C., Imielinski, T., September 1997. Geocast – geographic addressing and routing. In: Third ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'97). ACM, pp. 66–76.

Nittel, S., Duckham, M., Kulik, L., October 2004. Information dissemination in mobile ad-hoc geosensor networks. In: Third International Conference on Geographic Information Science (GIScience'04). Vol. 3234 of Lecture Notes in Computer Science. Springer, pp. 206–222.

Ostermaier, B., 2005. Analysis and improvement of inter vehicle communication security by simulation of attacks. Master's thesis, Technische Universität München, Institut für Informatik, Munich, Germany.

Piórkowski, M., Raya, M., Lugo, A. L., Papadimitratos, P., Grossglauser, M., Hubaux, J.-P., January 2008. TraNS: realistic joint traffic and network simulator for VANETs. ACM SIGMOBILE Mobile Computing and Communications Review 12 (1), 31–33.

Tuttlies, V., Schiele, G., Becker, C., November 2007. COMITY – conflict avoidance in pervasive computing environments. In: On the Move to Meaningful Internet Systems 2007: OTM 2007 Workshops. Vol. 4806 of Lecture Notes in Computer Science. Springer, pp. 763–772.

Wolfson, O., Sistla, A. P., Chamberlain, S., Yesha, Y., July 1999. Updating and querying databases that track mobile units. Distributed and Parallel Databases 7 (3), 257–287.

Wolfson, O., Xu, B., Sistla, A. P., June 2004. An economic model for resource exchange in mobile peer to peer networks. In: 16th International Conference on Scientific and Statistical Database Management (SSDBM'04). IEEE Computer Society, p. 235.

Wu, H., Fujimoto, R., Guensler, R., Hunter, M., October 2004. MDDV: a mobility-centric data dissemination algorithm for vehicular networks. In: First ACM International Workshop on Vehicular Ad Hoc Networks (VANET'04), in conjunction with MobiCom'04. ACM, pp. 47–56.

Xu, B., Ouksel, A. M., Wolfson, O., January 2004. Opportunistic resource exchange in inter-vehicle ad-hoc networks. In: Fifth International Conference on Mobile Data Management (MDM'04). IEEE Computer Society, pp. 4–12.

Young, K., Regan, M., Hammer, M., November 2003. Driver distraction: a review of the literature. Tech. Rep. 206, Monash University, Accident Research Centre.

Zhao, J., Cao, G., April 2006. VADD: Vehicle-assisted data delivery in vehicular ad hoc networks. In: 25th IEEE International Conference on Computer Communications (INFOCOM'06). IEEE Computer Society, pp. 1–12.

List of Table Captions

List of Figure Captions

1 Mobility and direction vectors 46

2 Illustration of a change in the frame of reference 47

3 Geometrical representation of Δd and Δt 48

4 Itineraries and events simulated 49

5 Percentage of vehicles informed 50

6 Benefits of adaptive vectors 51

7 Evolution of the messages exchanged in low traffic conditions 52

8 Evolution of the messages exchanged in high traffic conditions 53

9 Time needed to receive a message in low traffic conditions 54

10 Time needed to receive a message in high traffic conditions 55

List of Tables

List of Figures

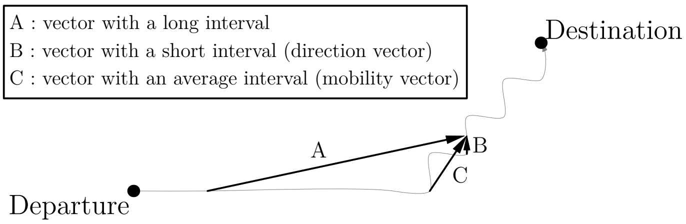

Fig. 1. Mobility and direction vectors

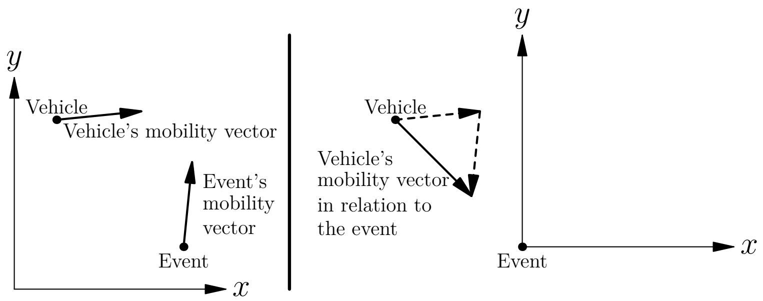

Fig. 2. Illustration of a change in the frame of reference

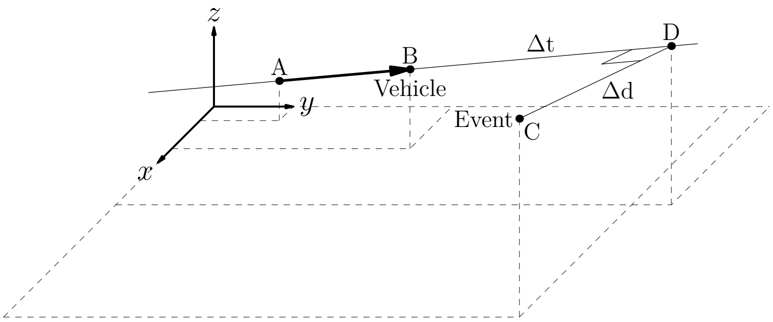

Fig. 3. Geometrical representation of Δd and Δt

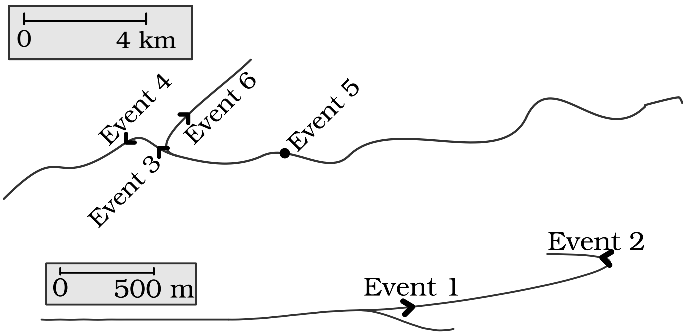

Fig. 4. Itineraries and events simulated

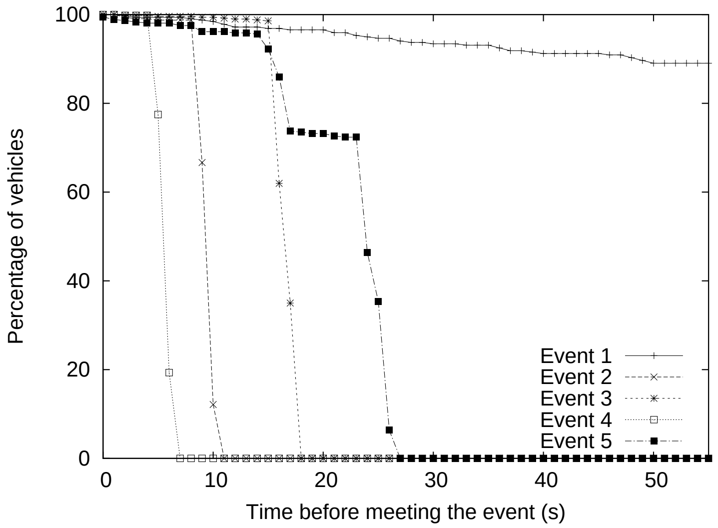

Fig. 5. Percentage of vehicles informed

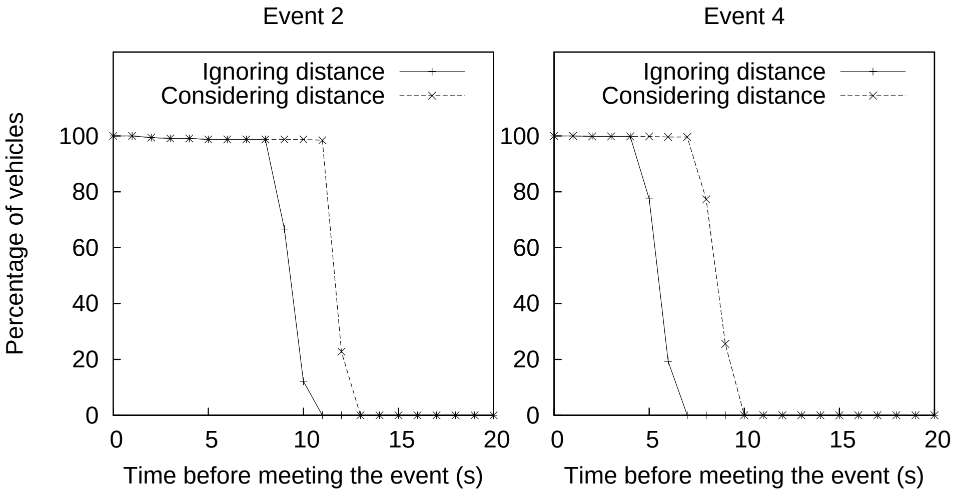

Fig. 6. Benefits of adaptive vectors

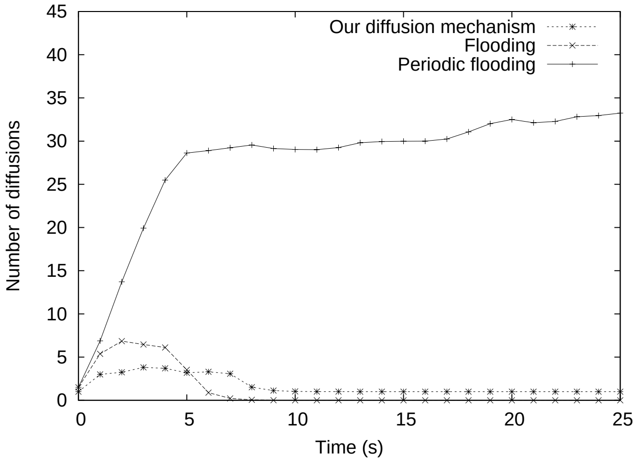

Fig. 7. Evolution of the messages exchanged in low traffic conditions

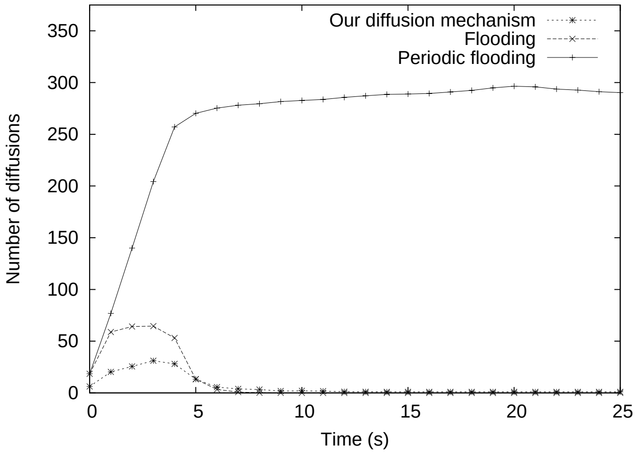

Fig. 8. Evolution of the messages exchanged in high traffic conditions

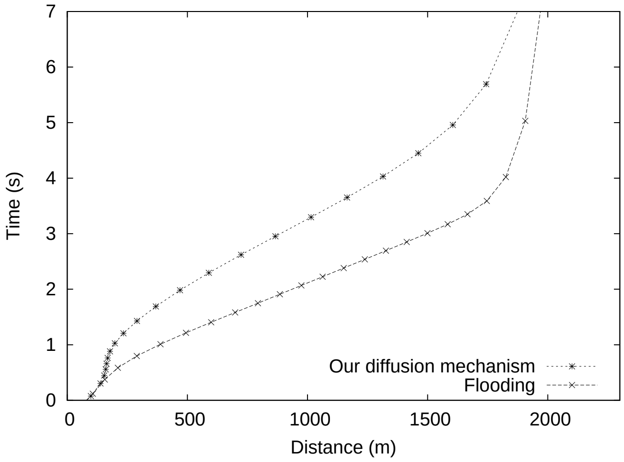

Fig. 9. Time needed to receive a message in low traffic conditions

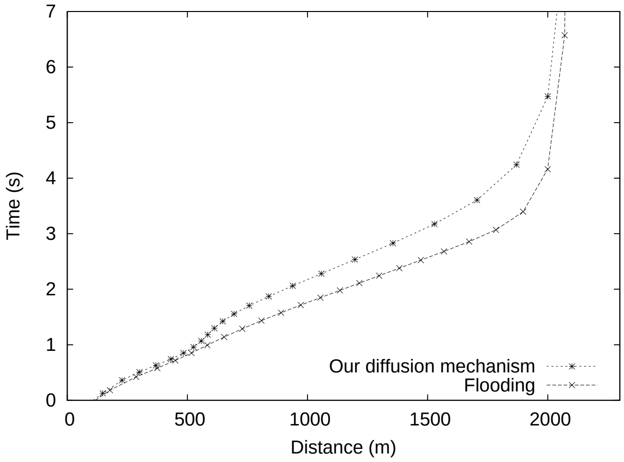

Fig. 10. Time needed to receive a message in high traffic conditions