Partage d'informations dans les réseaux ad hoc inter-véhicules avec VESPA

Nicolas Cenerario #1, Thierry Delot #2, Sergio Ilarri *3

#LAMIH Laboratory, University of Valenciennes
Le Mont Houy, 59313 Valenciennes - FRANCE
[1]`Nicolas.Cenerario@univ-valenciennes.fr`
[2]`Thierry.Delot@univ-valenciennes.fr`

**IIS Department, University of Zaragoza*
Maria de Luna 1, Zaragoza, 50018 - SPAIN
[3]`silarri@unizar.es`

***Résumé*— VESPA (Vehicular Event Sharing with a mobile P2P Architecture)[1] est un système permettant le partage d'informations décrivant divers événements dans un réseau ad hoc inter-véhicules (VANET). Contrairement aux autres systèmes existants dédiés à un type d'événement particulier, VESPA supporte tous les événements, qu'ils soient mobiles (e.g., véhicule dont le conducteur est en état d'hypovigilance) ou non (e.g., embouteillage), qu'ils concernent potentiellement tous les véhicules alentours (e.g., place de stationnement disponible) ou exclusivement ceux qui circulent dans une direction particulière (e.g., freinage d'urgence).**

Dans cet article, nous nous concentrons sur les fonctionnalités de base de VESPA. Nous décrivons notamment comment la pertinence des événements est évaluée grâce à une fonction appelée probabilité de rencontre. Ensuite, nous présentons le protocole de dissémination permettant la transmission des données décrivant les événements entre véhicules. Ce protocole limite le nombre de messages émis afin de ne pas saturer le réseau. Il adapte également la diffusion des événements aux véhicules potentiellement intéressés en fonction de leur type.

***Mots Clés*— VESPA, réseaux ad hoc inter-véhicules, événement, pertinence, probabilité de rencontre**

---

[1] Consultez la page: http://www.univ-valenciennes.fr/ROI/SID/tdelot/vespa/ pour plus d'informations.

I. INTRODUCTION

Today, 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. In spite of significant efforts to reduce the number of persons dying on the road, this number remains quite high, mainly due to the human factor (e.g., accident-prone behavior or low response time). 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.

Our work concentrates on data management in inter-vehicle ad hoc networks which rely on the use of short-range networks (about a hundred meters), like IEEE 802.11 or Ultra Wide Band (UWB) standards, for vehicles to communicate [12] and provide bandwidth in the range of Mbps. Using such communication networks, the driver of a car can receive information from its neighbors.

Many pieces of information can be exchanged in the context of inter-vehicle communications, for instance to warn drivers when a potentially dangerous event arises (an accident, an emergency braking, an obstacle on the road, etc.) or to try to assist them (with information about available parking spaces, traffic congestions, real-time traffic conditions on one road, etc.). Those different events can be detected by a car, using for example the numerous embedded sensors, and lead to the generation of a message transmitted to the other potentially interested vehicles, either directly or using multi-hop relaying techniques. Once received by a vehicle, the relevance of a message has to be evaluated, according to spatial and temporal criteria, in order to determine whether the driver should be warned or the message should be further broadcasted.

These last years, numerous research works have investigated how to disseminate data in inter-vehicle networks and how to evaluate their relevance. The different solutions proposed have addressed the exchange of one particular type of information and have proposed dedicated dissemination protocols. The originality of our solution resides in the ability to support any type of event (emergency braking, available parking spaces, etc.). Therefore, VESPA relies on the computation of an encounter probability, used to estimate whether an event is relevant for a vehicle (i.e., is the vehicle going to encounter the event?). The encounter probability is then used to adapt the dissemination according to the type of event.

The contributions of our work are summarized below:

- *Event classification* – we introduce a classification of the different events that may occur on the roads, taking every type of possible event into account in the same solution. Even mobile events, which are not considered in other works, 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.
- *Relevance evaluation* – Our technique is able to evaluate the relevance of the information exchanged between vehicles by defining a four-dimension mobility vector, computed for both vehicles and events. These vectors are then used to determine the probability of encounters between a vehicle and an event.
- *Dissemination protocol* – VESPA relies on a protocol able to disseminate different types of events according to spatio-temporal criteria, while avoiding network saturation.

The rest of the paper is organized as follows. Section II presents some related works. Section III introduces our classification of events and describes the representation of events exchanged with our VESPA system. Section IV focuses on the relevance of events and explains how the encounter probability is calculated. Section V presents our dissemination protocol. In Section VI, we discuss the experimental evaluations we have performed, both to validate the utility of the encounter probability and the dissemination protocol. Finally, Section VII offers our conclusions and gives some ideas for prospective research.

II. RELATED WORK

Inter-vehicle communication is a recent field of research. Nevertheless, some studies have already made significant contributions. In this section, we indicate some relevant related works.

In relation to network protocols, the FleetNet project (2000-2003) [10], [9], followed later by the NoW project[2] (finishing at the end of 2007), and CarTalk [14], [3] (2001-2004) worked to exploit inter-vehicle communications to make driving safer. Both FleetNet and CarTalk used multi-hop communication techniques, but while FleetNet was supported by a partial fixed infrastructure, CarTalk used no existing infrastructure. Those are very interesting projects which focus mainly on the network level of inter-vehicle communication. They use Geocast communication protocols [5], [13], which allow to determine the geographical area where messages have to be conveyed, either through direct inter-vehicle communication or multi-hop communication.

Recently, research has focused on data management and data dissemination rather than on specific network protocols for inter-vehicle communication. For example, we would like to highlight:

- In the context of the *TrafficView* system [15], [16], the authors consider vehicles moving on roads with multiple lanes on each direction. They present different dissemination protocols for periodically broadcasting all the data stored in a vehicle using a single network packet, thanks to data aggregation techniques. In [1], [2], the authors study the way to determine the *dissemination area*, which is the area where the data should be broadcasted.
- In [11], the authors explain that the network connectivity is a limiting factor for information dissemination since chains of vehicles are needed for broadcasting and a low traffic density may become a problem. The authors so make a clear distinction between data transport via locomotion (vehicles' movements) and via wireless communications to manage such situations. This distinction is also considered in the context of multi-hop routing protocols [20] with a *carry and forward* strategy, which aims at keeping the information stored on the car until it can be transmitted to another one. In the same category of solutions, we would also like to mention [18], which exploits the mobility of vehicles to perform the data dissemination.
- In the *Mobi-Dik* project [19], the techniques used to disseminate data do not require multi-hop communications. They are in fact much closer to 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 whether it should be stored in the cache and/or broadcasted later on.
- In [17], the authors focus on how to disseminate relevant geospatial information (mainly road hazards). They propose a dissemination strategy that they say to be similar to the opportunistic exchange proposed in Mobi-Dik. In addition, they introduce interesting metrics for efficiency such as *ignorance* (not knowing about a hazard found by the vehicle) or *redundancy* (receiving irrelevant information).

Summing up, the main goal of existing V2V communication solutions is to limit the number of messages exchanged to avoid overloading the network, which is indeed crucial if the correct functioning of the applications is to be guaranteed. The existing protocols and dissemination techniques mentioned above are generally interesting. Nevertheless, they only 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. Existing solutions are all specific to a particular type of event, presented as an application example, and thus cannot support other event types. For instance, Mobi-Dik provides a very interesting solution for the problem of information-sharing inside a restricted spatio-temporal area. Although these techniques are very well adapted for sharing information between cars about available parking spaces, they cannot be exploited to relay information about an accident or

---

[2] http://www.network-on-wheels.de/

an emergency braking situation. Thus, it is really difficult to compare existing solutions since they are well adapted to a particular type of event but cannot support other types of events. Besides, it is really difficult to imagine embedded systems in cars dealing with only one type of event and not with the others. Indeed, numerous events have to be taken into consideration to develop a new generation of driver assistance systems. Our interest is to be able to use a single data sharing approach valid for all types of events, in order to deploy a generic system into cars.

III. EVENTS AND DATA REPRESENTATION

To date, existing V2V solutions have considered only a small subset of the possible types of events, primarily focusing on stationary events. However, numerous types of events –both mobile and stationary– are possible, since there is a lot of information that drivers may find relevant, about: accidents, traffic congestion, emergency braking situations, fuel prices, available parking spaces, emergency vehicles such as ambulances, obstacles on the road, or the behavior of drivers (e.g., strange maneuvers due to intoxication or lack of vigilance[3]), to name but a few possibilities. In order to determine the relevance of events, it is first necessary to classify the different types of events. In the rest of this section, we propose a system of event classification and describe how these events are represented in our work. For simplicity, not only all kind of events but also road hazards and available resources are called *events* in the following.

A. Event classification

The solution that we propose in this paper not only supports stationary events, such as the presence of gas stations, but also mobile events, such as an emergency vehicle asking preceding vehicles to yield the right of way. When supporting such mobile events, the set of vehicles for which the event information is relevant evolves according to both the movements of the mobile event (in the example, the emergency vehicle) and the other vehicles involved (in the example, the preceding vehicles). None of the solutions presented in Section II supports mobile events. Besides, the direction of traffic is also of major importance in establishing the relevance of shared information, even for non-mobile events (e.g., consider a traffic congestion affecting only the vehicles moving in one direction).

So, we classify inter-vehicle network events in four different categories:

1) *stationary, non-direction-dependent events*;
2) *stationary, direction-dependent events*;
3) *mobile, non-direction-dependent events*;
4) *mobile, direction-dependent events.*

By *direction-dependent events* we mean events that are not relevant for all nearby vehicles, but only for the vehicles traveling in a particular direction. On the other hand, *mobile events* are (as explained before) events whose locations change along time.

Let us illustrate our classification system by giving 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, regardless of the 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. The vehicles close to the accident but moving in the opposite traffic stream should ignore the message so as not to distract the driver and cause a second accident. Messages warning vehicles of the lack of vigilance of a person driving on a two-way road is a mobile, non-direction-dependent event because it concerns 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. Our goal in proposing such a classification of events is to support, in the same solution, all the types of events which can occur on the roads.

B. Data Representation

In our solution, the four types of events identified in the previous section are used to represent all events occurring on the roads. In the following, we describe how these different events are represented

---

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

when created[4] in order for them to be exchanged between vehicles (a summary of the attributes considered is shown in Table I):

Attribute Name	Type
Key	string
Version	int
Importance	int
CurrentPosition	PositionAndTime
DirectionRefPosition	PositionAndTime
MobilityRefPosition	PositionAndTime
LastDiffuserPosition	PositionAndTime
HopNumber	int
Description	EventDescription

TABLE I
EVENT REPRESENTATION

- Each event is first characterized by a unique *Key*, generated by concatenating a unique vehicle identifier (for example, its MAC address) with a local unique event identifier.
- A *Version* number is also attached to each event 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 to other vehicles, for example because a mobile event has moved, the vehicle which created the event may produce a new version of the same event.
- An *Importance* is associated to each event. This attribute helps to determine whether the information should be presented to the driver or not. Thus, to avoid disturbing drivers, not all the relevant information should be presented to them. For example, if a vehicle approaches an available parking space, the driver is informed only if such information has been requested.
- The *CurrentPosition* attribute indicates the time and place corresponding to the generation of the event.
- Two different preceding reference positions and their timestamps (*DirectionRefPosition* and *MobilityRefPosition*) are also stored. These markers allow each vehicle to receive information to evaluate the mobility and direction of an event (see Section IV-A), which is necessary in order to estimate the event's relevance.
- The *LastDiffuserPosition* contains the position of the last vehicle which relayed the message. It is used by the dissemination protocol (see Section V).
- The *HopNumber* attribute indicates the number of broadcasts of the message (see Section V).
- Finally, a *Description* field describes more precisely the represented event (e.g., accident, emergency braking, etc.). This field is used to transmit concrete information to drivers when they need to be warned.

Let us note that the positions shown in Table I correspond to GPS statements and thus include 3-dimensional coordinates as well as a statement of the GPS time. Using the GPS time allows us to avoid synchronization problems between the internal clock of the vehicles.

Notice that the type of the event (stationary or mobile, direction-dependent or not) is not explicitly represented as an attribute of the event. Indeed, the type of event can be easily deduced using some of the other message fields. Moreover, we will show in Section IV that each type of event is managed in the same way when received on a vehicle. On the contrary, when an event is generated, its type has to be known in order to fill the right attributes before transmitting the information to other vehicles. More precisely, the values of the reference positions (*MobilityRefPosition* and *DirectionRefPosition* attributes) depend on the type of event considered. In fact, these positions correspond to former positions of the event and will be used to estimate its direction and speed, as we will further explain in Section IV. When dealing with a stationary object/event, the *MobilityRefPosition* will always be equal to the value of the *CurrentPosition* attribute. Similarly, for non-direction-dependent events the value of *DirectionRefPosition* will be set to *null* to allow the identification of such type of event.

Several versions of the same event may be generated. Such versions allow to update the event, including its mobility profile. So, for example, in case the car of a sleepy driver is blocked due to a traffic congestion, new versions of the generated

---

[4] We will not consider Human Machine Interface (HMI) aspects in this article. We rather focus on the representation and relevance estimation of events. The creation of those events may be initiated by devices embedded in the vehicles (for example by coupling the airbag system with the creation of an event representing an accident).

mobile non-direction-dependent event (warning about his/her lack of vigilance) would indicate the speed reduction. If the car of that driver should stop for a while for the same reason, then the new versions of the event would include two reference positions equal to the current position and then represent a temporary stationary event.

IV. COMPUTING THE RELEVANCE OF EVENTS

One of the major problems in V2V communications is determining the relevance of an event to a receiving vehicle. In this section, we first show how the data describing an event is exploited to compute a four-dimensional mobility vector (three dimensions for the spatial coordinates and one for the temporal dimension). Then, we introduce the notion of encounter probability, used to estimate whether a vehicle is expected to encounter an event or not in order to decide if the driver should be informed about it. We have chosen not to use a navigation system or roadmap. Although this could improve the precision of our proposal, it also presents some difficulties (see Section IV-B). Thus, our main goal is to prove the usefulness of the encounter probability.

A. *Mobility & Direction Vectors*

To estimate the direction of a moving object, we use vectors that run between a preceding position (called the "reference position") and the object's current position. These vectors are used to situate vehicles as well as mobile events. First, the position of object A at time t is expressed as:

$$P_A(t) = \begin{pmatrix} x_{A_t} \\ y_{A_t} \\ z_{A_t} \\ t \end{pmatrix}$$

where x_{A_t}, y_{A_t} and z_{A_t} are the geographical coordinates of object A at time t. The mobility vector for object A between t_1 and t_2 is thus defined as:

$$V_A(t_1, t_2) = P_A(t_2) - P_A(t_1) = \begin{pmatrix} x_{A_{t_2}} - x_{A_{t_1}} \\ y_{A_{t_2}} - y_{A_{t_1}} \\ z_{A_{t_2}} - z_{A_{t_1}} \\ t_2 - t_1 \end{pmatrix}$$

Each vehicle is able to compute its own mobility vector. By applying this vector to the current vehicle position, an estimation of its future position is obtained:

$$P_A(t_{n+i}) = P_A(t_n) + V_A(t_{n-i}, t_n)$$

The estimated future position is highly dependent on the time interval selected between two position statements. Thus, if t_n and t_{n-i} are far away, the estimation of the future position is not precise but provides an overall impression of the object's direction. If the time interval is shorter, then the estimation is much more precise on the short term but no global view of the displacement can be observed. As an example, see arrows A and B in Figure 1.

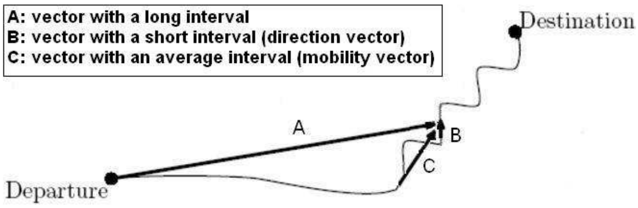

Fig. 1. Mobility and direction vectors

Depending on the way we select the time interval $[t_{n-i}, t_n]$, we distinguish:

- The *direction vector*, which is computed with a short interval (see arrow B in Figure 1). It provides a quite precise estimated future position but only in the very short term.
- The *mobility vector*, whose role is to provide an overall impression of the object's movement in addition to a good estimated future position. To achieve a good compromise between choosing a large and choosing a short time interval (arrows A and B in Figure 1), an "average" interval is used to compute the mobility vector (see arrow C in Figure 1).

Similarly, each vehicle can compute the *direction and mobility vectors of the events* it receives. For that purpose, it uses the data associated to the events, and more precisely the *CurrentPosition* attribute and either the *DirectionRefPosition* or the *MobilityRefPosition* attribute, respectively.

For each event, the *vehicle's mobility vector in relation to the event* is computed by changing the frame of reference (see Figure 2). The mobility vectors of one vehicle and one event are represented on the left side of the figure, and the mobility vector after the frame of reference has been changed is shown on the right side. With the change of reference, the computation of the encounter probability is simplified since a single vector needs to be managed for each couple <vehicle, event>, regardless of the type of event.

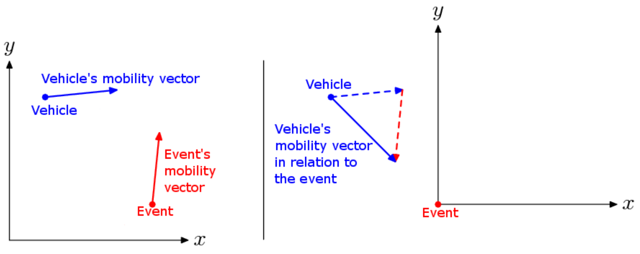

Fig. 2. Change in the frame of reference

To explain how this change of reference is computed, let us consider the mobility vectors of a vehicle A between t_{A_1} and t_{A_2} and of an event B between t_{B_1} and t_{B_2}:

$$V_A(t_{A_1}, t_{A_2}) = \begin{pmatrix} x_A \\ y_A \\ z_A \\ \Delta t_A \end{pmatrix}$$

$$V_B(t_{B_1}, t_{B_2}) = \begin{pmatrix} x_B \\ y_B \\ z_B \\ \Delta t_B \end{pmatrix}$$

The first step is to modify the mobility vectors in order to manage the same time basis (fourth dimension) for both vectors. We so obtain the following 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, we subtract the two vectors to obtain the mobility vector of vehicle A in relation to event B:

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

In the case of a stationary event, the mobility vector of the event is the *null* vector (whatever the time basis is):

$$V_{AB}(t_{A_1}, t_{A_2}, t_{B_1}, t_{B_2}) = \begin{pmatrix} x_A - 0 \\ y_A - 0 \\ z_A - 0 \\ \Delta t_A \end{pmatrix} = V_A(t_{A_1}, t_{A_2})$$

These vectors can then be used to compute an encounter probability to determine whether a vehicle will meet or not an event, as we describe in the following.

B. Encounter Probability

Using the mobility vector of the vehicle in relation to the event, the position of the vehicle, and the position of the event, we can deduce four elements which have an influence on the encounter probability:

- 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 event's generation time (stored in *CurrentPosition*) and the moment when the vehicle will be closest to it (Δg, *expected age of the event*).
- The angle between the direction vectors of the vehicle and the event (denoted by a colinearity coefficient c).

As an example, Figure 3 shows the geometrical representation of Δd and Δt in a certain scenario. To facilitate the graphical representation of the mobility vector in relation to the event, we consider a stationary event, but the principle would be the same with a mobile one. In the figure, B represents the vehicle position, C the event position, and $\overrightarrow{AB}$ is the mobility vector of the vehicle in relation to the event. Point D can then be determined, which allows a right-angled triangle to be constructed in D with $[BC]$ as hypotenuse. D is the closest point to

C on the straight line between A and B. $|DC|$ ($= \Delta d$) represents the minimal geographical distance between the vehicle and the event over time. $|BD|$ 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.

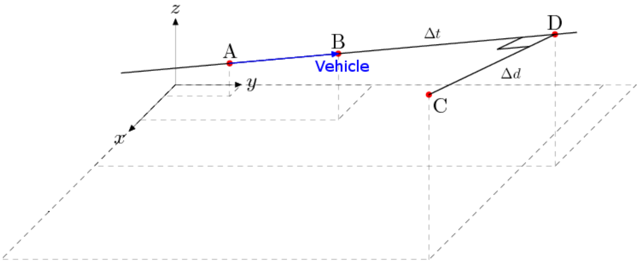

Fig. 3. Representation of Δd and Δt

As explained previously, the vehicle estimates its direction vector and the event's direction vector. From these two direction vectors, a *colinearity coefficient* (c) is obtained, which is a measure of the angle formed by the vectors. For direction-dependent events, this allows us to determine whether the directions of the vehicle and the event match. For non-direction-dependent events (identified because the *DirectionRefPosition* attribute is *null*), c is set to 0.

Once the Δd, Δt, Δg, and c values have been calculated, they are used to estimate an "encounter probability" between the vehicle and the event. The encounter probability (EP) is a value between 0% and 100%. It is computed using the following function:

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

where α, β, γ and ζ are penalty coefficients with values ≥ 0. They 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. For example, the greater the α value, the shorter the spatial range where the event is relevant. β and γ are used so that only the information about events that will be encountered very rapidly and the most recent information is considered. Finally, ζ is used to weigh the importance of the colinearity coefficient. Notice that if the vehicle is moving away from the event, then Δt is 0 and Δd is the current distance to the event. Therefore, the computation of the EP makes sense even in cases when an interesting event (e.g., a parking space) is behind us.

If a digital map is available and the driver's route is known, the values of Δd, Δt, Δg, and c could be computed with more precision (based on the map, the route, and a road-based distance measurement [6]). However, digital maps are not always available. Moreover, the intended route of the vehicle could be unknown. Even more importantly, in the case of mobile events, the expected route of the event will not be available, which would make the computation of the EP using the maps difficult.

The encounter probability allows to determine the relevance of an event. The greater its value, the more likely the vehicle is going to meet the event.

On the one hand, an event will be reported to a driver if the EP exceeds a given *relevance threshold* [8]. On the other hand, the EP is also used to adaptively disseminate the events [4], as we explain in the following.

V. DISSEMINATION PROTOCOL

We have primarily designed the encounter probability to determine whether an information contained in a message received on a vehicle should produce a warning for the driver. In this section, we do not consider the moment when the warnings have to be produced but rather want to demonstrate that the EP can also be used to disseminate information between vehicles. Our objective as concerns the dissemination protocol is to ensure that each vehicle for which an information is interesting will receive it. In the following, we first highlight the advantages of our dissemination solution, and then we present experimental results to evaluate these benefits.

A. Adaptive Dissemination Area

One of our main objectives is to disseminate different types of events (an accident, an emergency braking, an available parking slot, etc.) in the inter-vehicle network. Therefore, we have to support different dissemination modes. For example:

- An accident has to be diffused only to the vehicles driving on the affected lane.
- An emergency braking has to be diffused to the vehicles driving in a particular direction. Compared to the previous example, whereas the information about an accident should be relayed far away from the place where it took place, an event reporting an emergency braking should only be relayed a few hundred meters away.
- An available parking slot has to be transmitted to all close vehicles, whatever their direction, as it may interest them.

The use of the EP to determine the vehicles which have to broadcast the information they received allows to diffuse 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 in an area close enough to be relevant during the dissemination phase. For instance, the information about an available parking space would not be interesting for drivers several kilometers away from it. Therefore, in our dissemination solution, a vehicle computes the EP for each event it receives. If the value obtained is bigger than a certain *diffusion threshold*, it has to broadcast 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.

Besides, the EP also avoids the dissemination of obsolete events. The information diffused for the events evaluated in this paper (parking space, emergency braking, etc.) is only relevant for a short period of time. So, we focus our description on the way to reach the interested vehicles that must be informed about the event, rather than on the problem of trying to keep an information in the network during a given 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. Thus, our dissemination solution does not require any message to indicate the end of an event.

B. Limited Bandwidth Use

To avoid flooding, and so network congestions, our solution aims at desynchronizing the diffusions. Since the value of the EP may be greater than the diffusion threshold 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 broadcasting the message, although it will avoid rebroadcasting if it receives it again while waiting. The size of that period 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:

$$t = D \times (1 - \frac{d}{r})$$

where D is the maximum time to wait before broadcasting, r is the communication range of the wireless network used by the vehicles to communicate (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 stored 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 in advance which neighbor should rebroadcast would be unrealistic, as this would require that the vehicles track the locations of their neighbors

It should be noted that it may happen that none of the vehicles receiving a message about an event determines it relevant enough, after computation of the EP, to transmit 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 we use the following approach. Each time a vehicle diffuses a message, it waits during D seconds. Then, if it does not receive the message during that interval, it will periodically resend it until another vehicle estimates the event relevant and so diffuses it also. The *HopNumber* attribute (see Section III) is increased every time a vehicle relays a message, and it is used to determine if a message received is actually a retransmission. In that case, it should not be broadcasted again.

VI. EXPERIMENTAL EVALUATION

In this section, we present our experimental results. We first describe our experimental settings. Then, we evaluate the usefulness of the encounter probability as a metric to evaluate the relevance of the events. Finally, we evaluate our dissemination protocol.

A. Experimental Settings

For obvious scalability reasons, we evaluated our solution on a simulator rather than on the road. Choosing which simulator to use was a big challenge. It had to allow us to represent the environment, to define different routes for vehicles, and to integrate various speeds and traffic conditions, for example. When an event (whatever its type) occurred, we had to know which vehicles encountered that event and when. Finally, we had to be able to use all the available information to evaluate our dissemination and relevance algorithms. We studied the various simulators available (e.g., NS2[5], GloMoSim[6] and JiST-SWANS[7]), but it proved difficult to implement and evaluate our approach. In the end, we decided to develop our own simulator.

Our simulator allowed us to simulate realistic contexts. To represent different kinds of curves, the roads are represented using Bezier curves. This point is of major importance for us since we wanted to validate the use of mobility vectors on different kinds of roads (with different curve profiles). Each vehicle created on the simulator has its own position and speed, which are computed thanks to a mathematical function. 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 whether they should inform the driver or broadcast the message according to the value of the encounter probability.

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 us a more precise view whereas the second one corresponds to a part of the route between the cities of Valenciennes and Lille (in France). We considered six different events, located in strategic places, to evaluate our solution. All the events are stationary and direction-dependent, except *event5*, which is stationary and non-direction-dependent. Between 500 and 600 vehicles, with random trajectories, are simulated in each experiment. The speed of the vehicles usually varies between 90 km/h and 110 km/h, depending on the road (e.g., vehicles slow down in a curve).

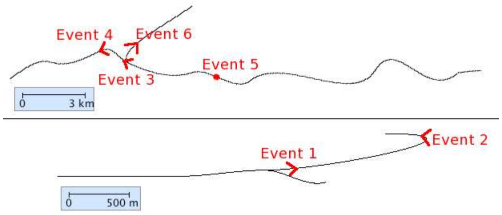

Fig. 4. Itineraries and events simulated

During our experimentations, we have used the following values for the penalty coefficients: $\alpha = 0.0033$, $\beta = 0.0010$, $\gamma = 10^{-8}$ and $\zeta = 0.25$. These values provide good results whatever the road configuration is, and so we will use them in the following to evaluate the interest of the EP. Naturally, it is possible to refine these coefficients according to the type of event (or even for the same event according to the context) in order to improve the results (e.g., for events representing parking spaces in urban areas, the penalty relative to the age of the event should not be the same on Saturday afternoon as on Monday). Similarly, based on experiments, we decided to compute mobility

---

[5] http://www.isi.edu/nsnam/ns/

[6] http://pcl.cs.ucla.edu/projects/glomosim/

[7] http://jist.ece.cornell.edu/

vectors using position statements performed every 500 meters, and direction vectors using position statements measured every 30 meters. The relevance/diffusion threshold is set to 75%.

B. *Evaluation of the Benefits of the EP to Estimate the Relevance of Events*

Our objective with our first set of simulations was to evaluate the concept of encounter probability. To only evaluate the efficiency of the encounter probability (and not the one of our dissemination protocol), we assume in this section that all the events are diffused using a flooding-based strategy. This assumption does not affect our goal, which is to prove the usefulness of the encounter probability to report the drivers about interesting events. The use of such a flooding-based technique to disseminate messages would not be appropriate for a deployment on real vehicles. Nevertheless, it allows us to separate the evaluation of the usefulness of the encounter probability from the evaluation of the proposed dissemination protocol.

In the simulations, we measured the percentage of vehicles which estimated an event relevant before meeting it (i.e., the encounter probability computed by these vehicles for the event is greater than 75%). The vehicles that do not meet the event are not considered in the computation of this percentage. In Figure 5, we present the evolution of warned vehicles as a function of the reaction time available before meeting the event (no vehicle travels on the road of *event6*, and so it does not appear in the figure since –as desirable– it is never reported). We draw the following conclusions:

- 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% of them had transmitted that warning 30 seconds before meeting the event. In that configuration, many vehicles generated a warning even if they never met the event. However, this was expected. Since this event is really close to the intersection of two roads, all 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 (in about 2 seconds).

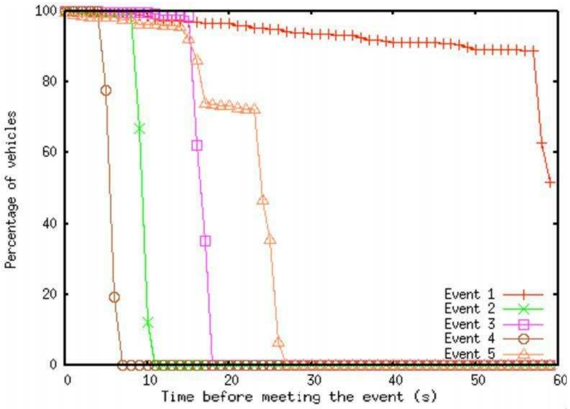

Fig. 5. Percentage of vehicles informed

- *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 the event *event3* confirm the good results obtained for *event1*.
- The results obtained for *event4* are not as good, with a warning generated only 4 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 encounter probability computed by the vehicles reached the threshold of 75% very late. Fortunately, this situation improves 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 us 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 8 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). We can conclude that adaptive vectors can be very useful in situations with strong changes of direction.
- *Event5* is considered relevant earlier than

event2, *event3*, and *event4*, as the travel direction does not penalize in the case of non-direction-dependent events.

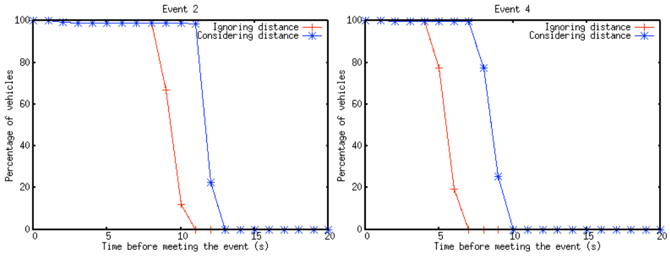

Fig. 6. Benefits of adaptive vectors

To summarize, the results obtained are rather satisfactory: the vehicles presented a warning to the driver in time, whatever the road configuration.

C. Evaluation of the Dissemination Protocol

In the following, we present the results of the simulations we performed to evaluate our dissemination protocol based on the encounter probability. Thanks to the results of our simulations, we could first observe that all the vehicles received the relevant events before meeting them, whatever the position of the event (straight line, curve, etc.). In the following, we focus on the evaluation of the efficiency of the dissemination protocol. Thus, our goal is to evaluate the number of messages emitted using our solution, to ensure that a message is not lost during its dissemination even if there are not many vehicles driving. Moreover, we also want to evaluate the time needed to deliver a message to a vehicle with our approach. In the following, we compare our solution with two others. The first one is based on a flooding technique. The second one performs a periodic flooding, motivated by the fact that some messages are lost with a traditional flooding when the traffic density is low (as we will indicate later). With periodic flooding, the messages are so re-emitted periodically.

We consider a single event and evaluate the different aforementioned dissemination strategies. In Figure 7, we present the total number of messages diffused at each second since the generation of the event in a scenario with low traffic density (about 1 vehicle every 100 meters). On the contrary, in Figure 8 we consider a scenario with high traffic density (about 1 vehicle every 10 meters). In the figures, we can observe the following:

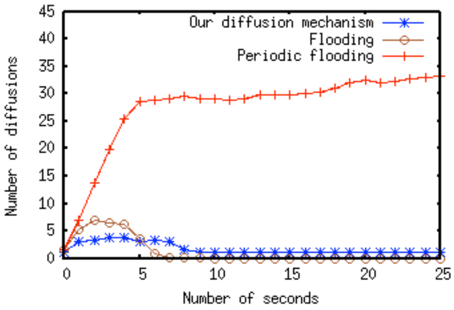

Fig. 7. Evolution of the messages exchanged in low traffic conditions

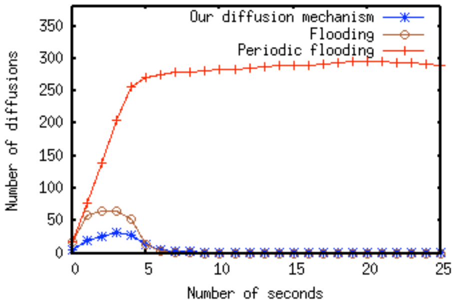

Fig. 8. Evolution of the messages exchanged in high traffic conditions

- Even with only one event, our dissemination solution strongly limits the number of messages exchanged.
- With our dissemination mechanism and traditional flooding, the number of messages decreases after a few 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).
- With the approach based on traditional flooding, the number of messages reaches 0 after a few seconds, and so the information about the event stops propagating. Therefore, a periodic flooding would be required instead.
- With our dissemination approach, the broadcast of the message continues at the tail of the diffusion chain[8] as long as the EP is high enough. Therefore, the previous problem is solved by our dissemination protocol.

---

[8] Even if it does not appear clearly in Figure 8 due to the scale, a few messages are still diffused after 10 seconds.

Since our solution may introduce waiting times at each hop of the dissemination process, we also wanted 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 our dissemination mechanism, the vehicles receive the information about the event a little more late. 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 her/him.

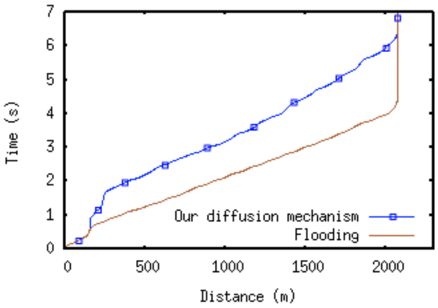

Fig. 9. Time needed to receive a message in low traffic conditions

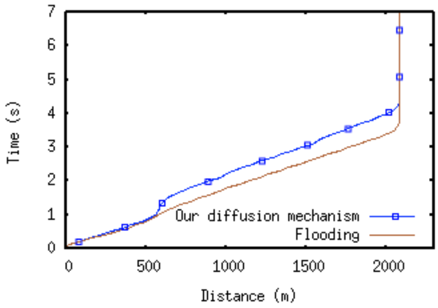

Fig. 10. Time needed to receive a message in high traffic conditions

VII. CONCLUSION & PERSPECTIVES

In this paper, we have presented a unified approach for disseminating data about different types of events in a vehicle network. Our proposal is based on the concept of encounter probability, which is computed to estimate the relevance of the events. As far as we know, this is the first proposal that does not focus only on a particular type of event. Moreover, our experimental results are really promising. On the one hand, we have shown the usefulness of the encounter probability to estimate the relevance of an event. On the other hand, some 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.

Currently, we are investigating ways to exploit data exchanged between vehicles to produce knowledge that can be used later on by drivers [7]. For example, it becomes so possible to dynamically detect potentially dangerous road segments or to determine the areas where the probability to find an available parking space is high.

ACKNOWLEDGEMENTS

This work was partly supported by the Nord-Pas-de-Calais region in the context of the D4S project. We also thank the support of a grant by the CAI-Europa XXI program.

REFERENCES

[1] C. Adler. *Information Dissemination in Vehicular Ad Hoc Networks*. PhD thesis, University of Munchen, 2006.

[2] C. Adler and M. Strassberger. Putting together the pieces - a comprehensive view on cooperative local danger warning. In *13th IST World Congress on Intelligent Transport Systems and Services*, 2006.

[3] D. D. Bruin, J. Kroon, R. Klaveren, and M. Nelisse. Design and test of a cooperative adaptive cruise control system. In *Intelligent Vehicles Symposium (IV'04)*, 2004.

[4] N. Cenerario, T. Delot, and S. Ilarri. Dissemination of information in inter-vehicle ad hoc networks. In *IEEE Intelligent Vehicles Symposium (IV'08), Eindhoven (The Netherlands)*. IEEE Computer Society, June 2008.

[5] A. Cheng and K. Rajan. A digital map/GPS based routing and addressing scheme for wireless ad hoc networks. In *IEEE Intelligent Vehicle Symposium (IV'03)*, 2003.

[6] A. Civilis and S. Pakalnis. Techniques for efficient road-network-based tracking of moving objects. *IEEE Transactions on Knowledge and Data Engineering*, 17(5), May 2005.

[7] B. Defude, T. Delot, S. Ilarri, N. Cenerario, and J.-L. Zechinelli. Data aggregation in VANETs: the VESPA approach. In *MOBIQUITOUS workshop on Computational Transportation Science (IWCTS'08)*. ICST, July 2008.

[8] T. Delot, N. Cenerario, and S. Ilarri. Estimating the relevance of information in inter-vehicle ad hoc networks. In *MDM International Workshop on Sensor Network Technologies for Information Explosion Era (SeNTIE'08), Beijing (China)*. IEEE Computer Society, April 2008.

[9] A. Festag, H. Füßler, H. Hartenstein, A. Sarma, and R. Schmitz. Fleetnet: Bringing car-to-car communication into the realworld. In *World Congress on Intelligent Transport Systems (ITS)*, 2002.

[10] C. Lochert, H. Hartenstein, J. Tian, H. Füßler, D. Hermann, and M. Mauve. A routing strategy for vehicular ad hoc networks in city environments. In *IEEE Intelligent Vehicles Symposium*, 2003.

[11] C. Lochert, B. Scheuermann, M. Caliskan, and M. Mauve. The Feasibility of Information Dissemination in Vehicular Ad-Hoc Networks. In *4th Conf. on Wireless on demand Network Systems and Services*, 2007.

[12] J. Luo and J.-P. Hubaux. A survey of research in inter-vehicle communications. In *Embedded Security in Cars - Securing Current and Future Automotive IT Applications*. Springer-Verlag, 2005.

[13] C. Maihofer. A survey of geocast routing protocols. *IEEE Communications Surveys & Tutorials*, 6(2):32–42, 2004.

[14] P. Morsink, R. Hallouzi, I. Dagli, C. Cseh, L. Schafers, M. Nelisse, and D. D. Bruin. Cartalk 2000: Development of a cooperative ADAS based on vehicle-to-vehicle communication. In *10th World Congress on Intelligent Transport Systems and Services*, 2003.

[15] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. 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, 2004.

[16] T. Nadeem, P. Schankar, and L. Iftode. A comparative study of data dissemination models for VANETs. In *3rd Annual Int. Conf. on Mobile and Ubiquitous Systems*, pages 1–10, 2006.

[17] S. Nittel, M. Duckham, and L. Kulik. Information dissemination in mobile ad-hoc geosensor networks. In *3rd Int. Conf. on Geographic Information Science*, 2004.

[18] H. Wu, R. Fujimoto, R. Guensler, and M. Hunter. MDDV: a mobility-centric data dissemination algorithm for vehicular networks. In *1st ACM Int. Workshop on Vehicular Ad Hoc Networks*, 2004.

[19] B. Xu, A. M. Ouksel, and O. Wolfson. Opportunistic resource exchange in inter-vehicle ad-hoc networks. In *5th Int. Conf. on Mobile Data Management*, 2004.

[20] J. Zhao and G. Cao. VADD: Vehicle-assisted data delivery in vehicular ad hoc networks. *IEEE Transactions on Vehicular Networks*, To appear, 2008.