Using V2V Communication to Create Over-the-horizon Awareness in Multiple-lane Highway Scenarios

Ramon S. Schwartz, Martijn van Eenennaam, Georgios Karagiannis, Geert Heijenk,
Wouter Klein Wolterink and Hans Scholten

Abstract—This paper describes the Over-the-Horizon Awareness (OTHA) protocol which provides an extended view of the traffic ahead to Driver Support Systems (DSS) by means of multi-hop vehicle-to-vehicle communication. The protocol extends and adapts the TrafficFilter [1] to make it suitable for multiple-lane highway scenarios. As basis, we rely on the information required by the Congestion Assistant [2], a driver support system that aids drivers in traffic jams. Simulation results show that OTHA achieves high node reachability and information accuracy.

Keywords: VANET, Traffic Jam, Driver Support Systems.

I. INTRODUCTION

Driver Support Systems (DSS) such as navigation systems and Adaptive Cruise Control (ACC) have been proposed to assist drivers on the road by improving safety, efficiency and comfort in the driving experience. In this context, a novel system referred to as the Congestion Assistant was proposed in [2] by Van Driel to assist drivers in the detection and traversal of traffic jams. Notably, such system can benefit and provide an enhanced experience to drivers when extended information of the traffic ahead is given as input. Acquiring this extended *view* comprises: (i) gathering and (ii) disseminating traffic information to multiple vehicles. Since large messages together with high transmission power and high transmission rates have shown to be the main reasons for radio channel congestion [3], the former action is tackled by aggregation mechanisms, e.g., in [4] or [5]. For the latter action, the broadcasting of messages is often performed and efficient mechanisms must be employed to cope with the *Broadcast Storm Problem* [6] in dense networks.

Based on the information required by the Congestion Assistant, Van Eenennaam introduced in [1], [7] the TrafficFilter which provides an *over-the-horizon awareness* of traffic jams ahead on the road. To the best of our knowledge, this is the only protocol that specifically copes with both aspects outlined above. However, the protocol was designed with single-lane straight highway scenarios in mind which raise the need for extensions and/or adaptations to make this solution adequate for multiple-lane scenarios.

This work is supported by iLAND Project, ARTEMIS Joint Undertaking Call for proposals ARTEMIS-2008-1, Project contract no. 100026; and by the Dutch Senter Novem/HTAS (High Tech Automotive Systems) Project Connect&Drive, Project no. HTASD08002

Authors are with Faculty of Electrical Engineering, Mathematics and Computer Science, University of Twente, 7500 AE Enschede, The Netherlands. Email: {r.desouzaschwartz, e.m.vaneennaam, g.karagiannis, geert.heijenk, w.kleinwolterink, hans.scholten}@utwente.nl

The contribution of this paper is that we introduce a novel multi-hop vehicle-to-vehicle communication protocol, the Over-The-Horizon Awareness (OTHA), that includes necessary extensions and modifications to the TrafficFilter in order to provide over-the-horizon awareness in multiple-lane highway scenarios. We consider the following highway scenarios: single and multiple-lane roads, junctions, and roads with multiple (opposite) directions.

The remainder of this paper is organized as follows. Sec. II provides a brief overview of the TrafficFilter. Based on this overview, in Sec. III we derive requirements that enable OTHA to cope with the multiple-lane scenarios considered. In the sequel, Sec. IV describes OTHA in detail. Sec. V describes the performance evaluation of the protocol carried out by means of simulations. Finally, Sec. VI concludes this paper and outlines our future plans.

II. THE TRAFFICFILTER

The TrafficFilter [1] follows a simple approach to aggregate traffic information. By means of ad hoc communication, vehicles collaboratively build a so-called TrafficMap. The TrafficMap is built with entries that contain the speed and position values of a few vehicles on the road, constituting a *speed profile* of the road. Based on thresholds, vehicles add, average and remove entries (samples) in order to ensure that the TrafficMap contains accurate traffic information. The underlying idea is that the first vehicle of a cluster initiates the TrafficMap by adding an entry with its own speed and position values. The TrafficMap is then relayed to other vehicles behind (upstream). Upon receipt, vehicles distributively take decisions by means of the *Sensitivity* ϵ function on whether to add a new entry or just relay the current information, defined as follows:

$$v_{new} = v_{own} \text{ iff } |v_{previous} - v_{own}| > \epsilon \quad (1)$$

where a new entry containing the current speed (v_{own}) and position of the vehicle is added if the difference between the previously added speed $v_{previous}$ and v_{own} exceeds the static threshold ϵ . In this way, only the subset of vehicles with valuable information contribute to the TrafficMap, thereby allowing for data compression. Other vehicles can still improve the TrafficMap by performing an *averaging* of their speed with the previous speed value inserted. Finally, the *reduce TrafficMap* function removes redundancies and merges similar entries in order to limit the TrafficMap size.

In order to disseminate TrafficMap information to every vehicle upstream, messages are propagated by means of broadcast communication. In order to cope with the *broadcast storm problem* in dense networks, a *suppression* technique is employed [8]. The underlying idea behind these suppression mechanisms is that when a vehicle is about to rebroadcast a certain message, the farthest vehicle in the message direction will have the highest priority to rebroadcast. This could be either based on probability, time, or a mixture of both. The remaining vehicles in the vicinity which also received the message can cancel (*suppress*) their transmissions when they receive an echo of that message. The TrafficFilter utilizes a modification of the slotted 1-Persistence scheme proposed in [8] and is referred to as the microSlotted 1-Persistence [9]. This modification seeks to break the synchronization that occurs when nodes assigned to a common time slot rebroadcast simultaneously and cause collisions as we explain later in Sec. IV-B.

III. REQUIREMENTS TOWARDS MULTIPLE-LANE HIGHWAY SCENARIOS

TrafficFilter provides means for capturing a compressed overview of the traffic assuming a single-lane road. However, for multiple-lane scenarios it lacks: (i) support for multiple lanes and capture of their individual behavior; (ii) support for the existence of multiple TrafficMaps originated in different roads; and (iii) coordination of message exchange in heterogeneous environments with vehicles driving in opposite directions and multiple lanes.

Internal Information	External Information
Own speed	Speed of tail of the jam
Own position	Speed of head of the jam
Distance to vehicle in front	Distance to vehicle in front
Own lane number	Position of the tail of the jam
Road segment	Position of the head of the jam
Own Road Identification	Lane number(s) of the traffic jam
Own driving direction	Road identification of the traffic jam
Junction point location	Driving direction of the traffic jam
Type of junction point	

TABLE I

REQUIRED INFORMATION CONSIDERING THE ADDITIONAL SCENARIOS

In order to support feature (i) we must capture and uniquely characterize traffic information with respect to the exact location where each piece of information has been generated, namely, on which road, direction, and lane. Based on these new requirements, we show in Table I the list of required information previously derived in [7] combined with new required information to address multiple-lane scenarios (in bold face). We assume that part of the required information is obtained from the vehicle itself by means of *internal* sensors and another part is obtained from other vehicles (*external*) by means of multi-hop wireless communication. By acquiring this information, the TrafficMap can effectively identify and represent speed deviations for individual lanes of various roads and directions. Furthermore, it enables the Driver Support Systems such as the Congestion Assistant to aid vehicles traversing traffic jams, or propose alternative

routes by evaluating upcoming junction points' location and type (exit or entrance points). Since the traffic behavior of hundreds of kilometers away is not so relevant to drivers, we limit the area of awareness to a *road segment*. OTHA is responsible for acquiring the external information. Each entry containing the speed and position of a vehicle is accompanied by the time stamp of its inclusion and mapped to the corresponding road, direction and lane values. This can be accomplished by matching the position values to the map provided by a navigation system.

In order to support features (ii) and (iii) OTHA adapts and provides new functions for managing and disseminating TrafficMap information. Specifically, in the information managing process we adapt functions used in the TrafficFilter to support multiple lanes and introduce new functions such as *merging*, *Discard Non-relevant Information*, *Prepare TrafficMap Information* to cope with TrafficMap data received from different roads, directions and lanes. With regard to the dissemination process, we adopt a new suppression policy that prioritizes messages sent by vehicles that included new information to the TrafficMap. The goal is to increase data accuracy. These functions are described in the next section.

IV. THE OTHA PROTOCOL

In this section the Over-The-Horizon Awareness (OTHA) protocol is described. Its basic functioning is shown in Fig. 1 with an example. Each vertical line represents a vehicle moving at a certain speed. The decrease in speed in the first two lanes indicates a traffic jam while the almost constant speed in lane three indicates free-flow traffic. We denote as *TrafficMap flow* the flow of traffic information propagated to vehicles upstream by means of multiple TrafficMap Messages. A flow is initiated by the *flow initiator* vehicle, in this example positioned close to kilometer 10. The flow initiator builds the TrafficMap by including the TrafficMap identification headers. Upon the receipt of a new TrafficMap Message, other vehicles evaluate the previous speed included in the TrafficMap and compare it with their own speed following Eq. 1. Vehicles either add a new entry and are referred to as *source vehicles* or simply forward the current TrafficMap and are referred to as *relay vehicles*. Since flow initiators also include a first entry with their own information, they are a special class of source vehicles. The result of this process is that for a certain observer, vehicle at position zero, the TrafficMap contains fundamental entries with the speed and position of a few vehicles (points in the figure) that gives an overview of the current traffic condition ahead.

Fig. 1. The threshold-based approach for a road with multiple lanes

The protocol is divided into two layers: the *Traffic Filter Protocol Layer* which aims to manage the TrafficMap data, and below it the *Dissemination Protocol Layer* which allows for quick and efficient TrafficMap data dissemination.

A. The Traffic Filter Protocol Layer

The Traffic Filter Layer functioning is depicted in Fig. 2. The main functionality is the *TrafficMap Manager*, which is a set of functions responsible for maintaining and updating the TrafficMap stored in vehicles. This process is triggered by the following events: (i) upon the receipt of a new TrafficMap Message originated by other vehicles and received from the lower layer; (ii) by the occurrence of pre-determined internal events within the vehicle; (iii) or by a TrafficMap data request received from the lower layer.

Fig. 2. The Traffic Filter Protocol Layer

The first event is meant to address the rebroadcasting of TrafficMap information. This information arrives from the lower layer by means of the *Rebroadcasting path*. From this path, the information received by other vehicles is first analyzed and all “non-relevant” information is discarded by the *Discard Non-Relevant Entries* function. The decision on which information is relevant or not will depend on the application running above OTHA, e.g., the Congestion Assistant. For the sake of simplicity, we define that roads where the vehicle is not able to go within the current *road segment* are not relevant. After this first filtering of information, if there is still some relevant information left, it will be merged with the current stored TrafficMap information by means of the *Merging* function. This process will keep the most up-to-date information regarding each lane by means of time stamp values. This function is crucial to capture and merge multiple TrafficMap flows that, for instance, are originated in different roads separated by a junction down the road.

The following functions have been previously proposed in [1] and are briefly explained. The *reduce TrafficMap* function removes redundancies and keeps the current TrafficMap size below a certain limit. One novelty is the use of time stamps in the removal decision process and the higher priority given to entries regarding congested areas of the road, since they are of higher importance to DSS systems. Following the protocol, the *sensitivity* ϵ decides whether a new entry must be added to the TrafficMap based on the last entry added to the lane on which the vehicle is currently situated, as explained in Sec. II. New entries are added by the *Add Sample* function. Whenever the *sensitivity* ϵ decides not to add a new entry, the vehicle can still improve its TrafficMap by performing an averaging of its current speed with the last speed value received for its own lane by means of the *averaging* function. This averaging is only performed up to the static threshold distance Δ from the vehicle which added the previous entry (position $P_{own} - P_{prev}$), since very distant vehicles may not be representative for that entry anymore. Finally, a request is sent to the lower layer in order to disseminate the updated TrafficMap to other vehicles. Moreover, whenever the process reaches this point, the current TrafficMap information is sent up to the application. As part of a rebroadcasting process, other vehicles upstream must also receive the updated TrafficMap even if no entry has been added. Due to this fact, the last decision step of the TrafficMap Manager will always allow the sending of the mentioned request to the lower layer. As there are decision processes in the lower layer that rely on the current vehicle type, the *message request* includes information about whether the vehicle is a source vehicle, i.e., it has added a new entry to the TrafficMap, or is simply a relay vehicle, i.e., it has simply performed the averaging process.

The second form of initializing the TrafficMap Manager process concerns events that are triggered by the *Internal Event Manager* by means of the *Event Path*. In this path, the functions executed are basically the same when compared with the process initiated by the *Rebroadcasting path*. The exceptions are the *merging* and *discard non-relevant entries* functions, since now no information from other vehicles is received and, therefore, the execution of these functions is not necessary. Because the vehicle is not participating in a rebroadcast process, the last step in the TrafficMap Manager process will only permit a request to be sent to the lower layer for dissemination of the current information in situations when a new entry has been added to the TrafficMap. Otherwise, the process is simply finished.

In this work, the following internal events are considered:

- **The periodical speed check/reduce TrafficMap timer has expired:** it forces the vehicle to compare its current speed with the last speed value added to the TrafficMap for its current lane. When a sudden speed change occurs, the *sensitivity* ϵ function will allow a new entry to be included in the TrafficMap and a new message will be requested to warn vehicles behind about it. As part of event path, the *reduce TrafficMap function* guarantees that old and redundant entries are periodically removed.

- **The vehicle has moved to another lane:** whenever a vehicle moves to a different lane, its current speed might deviate considerably from the last speed value added to that lane. For such situations, a new entry will be added to the TrafficMap and a message request will be sent to the lower layer to warn vehicles upstream.

- **The vehicle has moved to a different road:** similar to a lane change. However, when moving to a different road the *Reduce TrafficMap* function can discard all previous entries concerning the previous road, since it might not be relevant.

The last form of triggering the TrafficMap Manager regards the receipt of a data request from the lower layer. The *Prepare TrafficMap Information* process ensures that the most up-to-date information is included in the TrafficMap Message just before it is sent to other vehicles. When the lower layer defines the vehicle as *flow initiator*, identification headers (road ID, road direction, and time stamp) will be included in the TrafficMap. In addition, all entries regarding the current vehicle's road are erased and a new entry is added to the TrafficMap. Finally, the whole data is retrieved and sent back to the lower layer. Vehicles which are not *flow initiators* simply include the most up-to-date TrafficMap information and return it to the lower layer.

B. The Dissemination Protocol Layer

Current broadcast suppression schemes such as the ones described in [8] have been proposed as general solutions for vehicular communication. When used in the task of disseminating traffic information to derive a speed profile of the traffic ahead, these solutions need to be tailored. As explained previously, two types of vehicles are considered: *source* and *relay* vehicles. When only relay vehicles are considered any broadcast suppression technique suffices, since the goal is to simply relay messages with minimum end-to-end delay. However, *source* vehicles contain crucial information to be included in the TrafficMap that must always be considered. Consequently, in our dissemination protocol we do not allow source vehicles' broadcasts to be suppressed. This is guaranteed by assigning unique IDs to source vehicle messages. Suppressions occur in relay vehicles, as they repeat the previous ID received when rebroadcasting. One consequence of this design decision is that since messages will be broadcasted asynchronously by different source vehicles, a flow initiated by the head vehicle of a cluster might be split into multiple TrafficMap information flows along the road. On the one hand, multiple TrafficMap information flows result in a less time efficient protocol because of the higher number of messages transmitted. On the other hand, we clearly prioritize accuracy, since every speed deviation detected by a vehicle is broadcasted.

In this work, we base our design for the dissemination protocol layer on the TrafficFilter protocol. Accordingly, the suppression strategy is mainly based¹ on the *slotted 1-*

Persistence as it has shown to achieve the best performance among the techniques proposed in [8]. In order to cope with *source* vehicles, early time slots are reserved for them. Because they possess critical information, these early time slots give them the opportunity to transmit quickly and cancel TrafficMap Messages scheduled by relay vehicles.

Another important characteristic we consider is the speed dependency among vehicles. A reduction in speed by vehicles ahead on the road may induce a reduction in speed by vehicles coming behind. These are *events* that occur successively and towards vehicles upstream and are captured in our protocol by means of source vehicles. Thus, among source vehicles, the ones closest to the sender are given the earliest time slots. This measure aims at capturing and propagating events in the correct order. For relay vehicles we adopt the opposite pattern, as they do not possess new information and are meant to disseminate as quick as possible.

The time slot assignment for vehicles in the Dissemination Protocol Layer is defined as follows. When vehicle j further in message direction receives a message from vehicle i , it first calculates the percentage distance PD_{ij} between the two vehicles with respect to the *estimated* transmission range R .

$$PD_{ij} = \left\lceil \frac{\min(D_{ij}, R)}{R} \right\rceil \quad (2)$$

where D_{ij} is the relative distance between vehicles i and j . As a result, the PD_{ij} value will vary within the interval $(0,1]$ with large distances being closer to 1.

The time slot number S_{ij} assigned to either a source or relay vehicle j is defined by the following equation:

$$S_{ij} = \begin{cases} \lceil NS_{source} \times PD_{ij} \rceil - 1 & \text{if source;} \\ NS_{source} + \lceil NS_{relay} \times (1 - PD_{ij}) \rceil & \text{if relay.} \end{cases} \quad (3)$$

where NS_{source} and NS_{relay} are the total number of time slots reserved for source and relay vehicles, respectively.

Given time slot number S_{ij} , the total amount of time T_{Sij} vehicle j waits before rebroadcasting is given by equation:

$$T_{Sij} = S_{ij} \times st \quad (4)$$

where the slot time st is an estimation of the one-hop delay including the medium access delay and propagation delay.

The assignment of different time slots to vehicles depending of their positions clearly breaks the synchronization present in the simple flooding approach, where all nodes would simply rebroadcast simultaneously and cause collisions. However, a similar synchronization on a smaller scale can still occur with vehicles assigned to a common time slot. In [9] such issue has been identified and tackled with a variation of the slotted 1-Persistence technique called *microSlotted 1-Persistence Flooding* by introducing. This is achieved by staggering the wait time of the slotted scheme by means of *microslots*. These microslots have the duration of one DIFS in the 802.11p standard [10] and are allocated

¹A typo with regard to the ceiling math function position has been identified in the formula for the Slotted 1-Persistence technique proposed in [8], which leads to inaccurate distribution of vehicles among different time slots.

based on distance. Differently, the work described in [11] refers to this issue as the *Timeslot Boundary Synchronization Problem* and describes design guidelines for extra measures to be taken not only in the network layer but also in the link layer. This is achieved by inserting a pseudo-random delay to SIFS in the link layer. However, in congested networks an additional delay introduced only in the network layer does not suffice when nodes experience high contention in the link layer, as their timeslots could again align.

As in [11], we support the position that the synchronization must be broken in both network and link layers to be completely effective. However, as a preliminary solution we follow the guidelines proposed in [11] only for the network layer. In this way, we study the viability of this solution with the existing IEEE 802.11p MAC protocol layer. According to these guidelines, the extra delay must be chosen from a near continuous interval to break the alignment of timeslot boundaries. Following the idea adopted for the assignment of time slots, source vehicles closer to the sender receives smaller delay values and relay vehicles have the opposite pattern. When vehicle j receives a message from vehicle i , the additional delay AD_{ij} is defined as follows:

$$AD_{ij} = \begin{cases} D_{max} \times PD_{ij} & \text{if source;} \\ D_{max} \times (2 - PD_{ij}) & \text{if relay.} \end{cases} \quad (5)$$

where D_{max} is the maximum allowed delay for each type of vehicle. The result is that for each type of vehicle each time slot is stretched with an equal fraction of D_{max} . Moreover, the beginning of each time slot is shifted by the accumulated additional time of earlier time slots, thereby preserving the pre-defined st value and preventing overlapping between different time slots.

The time that vehicles have to wait before rebroadcasting is updated to include the additional delay described as follows in Eq. 6.

$$T_{Sij} = (S_{ij} \times st) + AD_{ij} \quad (6)$$

Our suppression strategy is exemplified in Fig. 3. Four time slots are utilized: the two earliest reserved for source vehicles and the remaining for relay vehicles. Source vehicles are marked with a rectangle surrounding each of them. The number above each vehicle indicates their turn according to their assigned time slot in Eq. 6. Among relay vehicles, the most distant ones from the sender are assigned to the earliest time slot: $t = 2 \times st$. Contrary, the closest source vehicles to the sender have the earliest time slot: $t = 0$.

Fig. 4 shows the Dissemination Layer process. A message is first received from the lower layer which is envisioned to be the MAC Layer defined by IEEE 802.11p, the upcoming IEEE standard for vehicular communication. The first decision process verifies whether the message has been originated by a vehicle *further in the message direction*. The message direction is included in each TrafficMap Message and is defined by the application running on top of OTHA. Here, we consider the direction to be against the traffic flow. The

Fig. 3. Overview of the dissemination protocol proposed

protocol starts as follows. If the message has been originated by a vehicle further in the message direction, the vehicle processing the message verifies whether the message is a rebroadcast of a *previous broadcast b*. If true and the vehicle has a *message scheduled for the same broadcast b as relay vehicle*, such message will be canceled (suppressed). As explained, messages from source vehicles are not canceled. If the message has been originated from a vehicle not further in the message direction, i.e., from a vehicle ahead on the road, a decision process evaluates whether that message has already been seen before. If it is an old message, it is simply discarded. Otherwise, the message is passed to the upper Traffic Filter Protocol Layer to be further analyzed.

Fig. 4. The Dissemination Protocol Layer

After the message is processed by the Traffic Filter Protocol Layer, a message request may be sent back to the Dissemination Protocol. When a message request arrives, it is first handled by the *Scheduler*. This process controls the sending of messages into the network. Three timers are defined: (i) the τ timer, (ii) Flood Free Period (FFP) timer, and (iii) the broadcast suppression timer, as proposed in [7].

The τ timer guarantees that a new TrafficMap is received by vehicles periodically. If no TrafficMap Message is received within this time, it means that the vehicle is the head of a cluster. Accordingly, it becomes the flow initiator and starts a new TrafficMap flow. The FFP simply ensures a minimum time between consecutive broadcasts to limit the level of radio congestion introduced by OTHA. In addition, the Scheduler also assigns the time slot of our suppression technique to each vehicle including the additional delay. The last step of the protocol is the preparation of the message to be sent down to the MAC layer by the *Message Builder*. This process is responsible for defining the message header and acquiring via a *data request* the latest TrafficMap data available in the vehicle's memory managed by the Traffic Filter Protocol Layer. Moreover, the data request includes information indicating whether or not the vehicle is the current flow initiator.

V. PERFORMANCE EVALUATION

In this section, we present the performance evaluation of OTHA carried out by means of simulations with OMNeT++ 4.0. Our goal is to verify whether OTHA functions properly under multiple-lane highway scenarios.

In our simulations, we utilize the Mobility Framework [12] and adjust the available implementation of the IEEE 802.11b protocol to comply with basic characteristics of 802.11p. In the MAC layer we set the bit rate to 6 Mbit/s (broadcast rate), the minimum Contention Window (CW) value to 15, slot time to 13 μ s, SIFS to 32 μ s, and DIFS to 58 μ s. In the physical layer, we run on the 5.87 GHz frequency band, with 10 MHz of bandwidth, and based on estimates we set the transmission power to 168.98 mW to achieve 500 m of interference range and 250 m of transmission range, assuming the Friis Free Space propagation model with pathloss exponent $\alpha = 3.5$.

In the Traffic Filter layer, we use the same values with which successful results have been obtained in [7]. Regarding the Dissemination Layer choices are made based on preliminary simulation results. Expiration times for τ and FFP timers are set to 3 and 0.1 seconds, respectively. The maximum additional delay D_{max} is set to 2.9 ms and the time slot st to 9 ms. We use 7 time slots where 5 are reserved for relay vehicles and 2 slots for source vehicles. Further study on traffic theory is required to determine the optimum values for the parameters utilized.

To evaluate the protocol we derive the following metrics:

- **Reachability**: the percentage of vehicles which receive each TrafficMap initiated. Ideally, dissemination protocols must achieve a percentage close to 100%.
- **Total Channel Utilization**: the percentage of busy time perceived by an arbitrary vehicle with respect to the total simulation time. The channel utilization takes into account transmitting time and any noise detected by a vehicle, i.e., errors or collisions. This metric evaluates how efficiently the medium is utilized by the protocol.
- **Delay**: the total time needed for a message to propagate from one end to the other on the road length considered in

each scenario. This is particularly important for critical information that must be disseminated as quickly as possible.

- **Accuracy**: is the focus of the protocol and is measured as the error (difference in speed values) of the data collected, i.e., the speed values of the entries added to the TrafficMap, compared with the real speed of vehicles present on the road. This value includes errors caused by: (i) the thresholds defined in the sensitivity ϵ function in the Traffic Filter Layer, (ii) propagation error of TrafficMap Messages in the vehicular network, and (iii) multi-hop propagation delay.
- **Distance of awareness**: in real mobility scenarios the connectivity among vehicles is not always certain. This metric serves then to measure, from the point of view of a fixed observer at one end of the road, the maximum distance of awareness it can achieve at an arbitrary time instant.

A. Static Scenarios

We define one static scenario for each of those considered in this work, as shown in Fig. 5. The arrow outside the road indicates the direction of vehicles on the road. The dark circle represents the *flow initiator* vehicle and the arrow next to it indicates the message direction of the TrafficMap information flow. All flow initiators are placed on the extreme side of each road towards the direction of vehicles. The vehicle represented with a white circle placed in the extreme opposite end of the road is used to gather relevant information for end-to-end delay and accuracy metrics and is referred to as *collector vehicle*. For each scenario, we perform 50 simulation runs for each of the following vehicle densities: 20, 40, 60, 80, and 100 vehicles/km/lane. The *flow initiator* vehicle creates the TrafficMap structure and starts transmitting at time instance zero. In each simulation run, 50 flows are initiated by this vehicle in intervals of 3 seconds.

Fig. 5. Illustration of the static scenarios considered

The distribution of vehicles along the road is determined by the Intelligent Driver Model (IDM) described in [13]. The IDM is a continuous car-following model which is essentially defined by an acceleration function. All scenarios illustrated are basically snapshots containing the speed and position of vehicles taken after the IDM model has been executed for a randomly chosen time. In every case, a traffic jam is induced by determining a lower maximum speed value in a certain region of the road, as observed for lanes 0 and 1 in Fig. 1.

Fig. 6. Performance Evaluation for Static and Mobility Scenarios. In order to show the statistical accuracy of the results, the plots above show the 95% confidence intervals. Figures (a) and (b) show results for static scenarios whereas (c) refers to results for the mobility scenario.

The performance of the protocol with respect to reachability is shown in Fig. 6(a). At density 20 vehicles/km/lane, the reachability is poor for almost every scenario due to a sparse network and consequent lack of connectivity among vehicles. Contrary, in the multiple-lane scenario the reachability achieves a high mark of almost 100%. This is explained by the existence of connectivity between vehicles in different lanes. As of density of 40, in the single-lane, junction, and opposite direction scenarios the reachability remains close to 100%. However, the reachability in the multiple-lane scenario decreases to 80% as at high densities. In fact, a similar behavior is observed in [9] for single-lane scenarios between densities 160 and 200, which corresponds to the number of vehicles within the transmission range for two lanes in densities greater than 80. The high number of vehicles within the transmission range increases the probability of collisions due to possible busy medium and simultaneous transmissions in the 802.11p MAC layer.

Fig. 6(b) shows the performance results for the multi-hop delay. At density 20, the multiple-lane scenario is the only one with a complete end-to-end connectivity among vehicles. The delay variation for the opposite direction, single-lane, and junction scenarios is similar, with a smooth decrease throughout the increase in density. A higher delay in low density situations is expected, since the time slots utilized by our broadcast suppression technique may not be equally distributed among vehicles, e.g., if there are only vehicles assigned to later time slots. The delay is generally higher in junction scenarios when compared, for instance, to single-lane scenarios due to existing multiple TrafficMap flows started in different roads. Since it is likely that these flows arrive upstream on the road asynchronously, later flows can be delayed by the FFP timer set by vehicles that rebroadcasted in early time slots during the travel of a previous flow. In the multiple-lane scenario, the high number of vehicles per time slot and the high probability of transmission collisions and errors result in a increase in delay in high densities.

In high densities, the channel utilization is perceived to increase. This is expected as the higher number of vehicles within a single time slot results in more transmissions and thus more receptions. Nevertheless, the channel utilization

remained less than 0.7% of the total simulation run time (vehicles were idle 99.3% of the time), which leaves room for other applications to run concurrently. This is explained by the fact that transmissions occur mainly in bursts, since TrafficMap flows are initiated at evenly spaced intervals determined by the τ timer expiration time, i.e., every 3 secs.

With regard to accuracy, simulation results show that the sampling error is limited by the low value of 1.1 km/h in sparse networks. As the density increases, the sampling error decreases, since vehicles at high densities drive at lower speeds due to the occurrence of traffic jams. Therefore, the speed deviation present on the road is also lower.

Overall, OTHA presents high reachability results for all scenarios with a small decrease perceived for dense multi-lane roads. This decrease is particularly expected when broadcasting with the 802.11p MAC protocol due to: (i) lack of acknowledgments, hidden terminal problem, and the constant small size for the Contention Window (16 slots) as it is never increased when broadcasting. In addition, the information is quickly delivered (below 0.6s) and accurate and yet without overloading the radio channel.

B. Mobility Scenarios

Fig. 7. Illustration of the mobility scenario considered

The mobility scenario combines variations of the addressed scenarios into one, as shown in Fig. 7. In this evaluation, we concentrate on the distance of awareness and accuracy metrics for evaluation. Accuracy implicitly evaluates delay, i.e., long delays result in old and inaccurate information, and the maximum distance of awareness gives indication of the reachability achieved. Results for channel

utilization have been analogous and are thus omitted. To ease the explanation of the results we divide the scenario into Sections 1, 2, 3 and 4. A single collector vehicle is placed in road R1 and there is one flow initiator in each road end. We perform 50 simulation runs, each run with a time duration of 300 seconds. Vehicles move at intervals of 0.5 seconds. As time evolves different and/or multiple vehicles are assigned as flow initiators as there are gaps and vehicles entering and leaving the road.

The distribution of vehicles is generated by means of the Quadstone Paramics 5.2 [14] traffic simulator executed with the CeeJazz plug-in. The generation rate of vehicles in the simulator is made high for Section 1, up to 40 vehicles/km/lane (illustrated in Fig. 7). After the junction point, the high density introduced in Section 1 is distributed among Sections 2 and 3, causing congestion. Due to an overall low density, we evaluate OTHA in a generally sparse scenario with vehicles moving at high speeds. Sections 1, 2, and 3 together account for over 80% of all traffic generated whereas traffic in Section 4 serves as radio background noise.

The distance of awareness achieved as time evolves is shown in Fig. 6(c). This plot illustrates the distance of awareness achieved (sampled) placed over the estimated maximum theoretical distance of awareness for Sections 1, 2, and 3. The maximum theoretical distance of awareness is calculated as the distance between the collector vehicle and the furthest vehicle downstream that could be reached with a transmission range of 250 meters (assuming no transmission errors). The sampled values are the distance averages achieved for all TrafficMaps received. The distance of awareness achieved in R1 (from Section 1 to 2) is in great part near the maximum theoretical distance achievable. From Sections 1 to 3, which includes R2 after the junction point, the plot shows some fluctuation and displacement especially at the beginning of the simulation. Such displacement is somewhat expected: (i) the maximum theoretical distance of awareness is calculated for each time instant and it does not take into account the end-to-end delay needed for the complete propagation of the TrafficMap. The time instance at which a TrafficMap is received may refer to an existing end-to-end connectivity a few seconds before, thus, it may be shifted to the right in this figure; (ii) Section 3 constantly presents a sparse network and therefore there is a lower probability that TrafficMaps are completely propagated in road R2; (iii) although a proper transmission power has been employed by vehicles to achieve 250 meters, the more distant vehicles are from each other the lower is the probability of successful communication.

With regard to accuracy, results indicate that the sampling representation error mean is 4.5 km/h with standard deviation close to 0.1 km/h. Considering the high speed variation of vehicles at some points in the mobility scenario, e.g., a sudden drop in speed from 100 down to 25 km/h in Section 1, a sampling error of 4.5 km/h is considerably low.

Overall, OTHA presents high distance of awareness during the whole simulation time. In addition, accuracy is high even in the presence of high speed fluctuations.

VI. CONCLUSION AND FUTURE WORK

In this work, we have proposed the OTHA protocol that comprises extensions and modifications to TrafficFilter [1], [9] in order to address multiple-lane highway scenarios in addition to single-lane straight highways previously addressed. In particular, we have addressed single and multiple-lane roads, junctions, and roads with multiple (opposite) directions. The performance of the OTHA protocol has been evaluated by means of simulations in both controlled environment with static scenarios and under a more realistic scenario consisting of vehicle traces with high mobility and speed variations. Our results show that OTHA preserves the high reachability achieved by TrafficFilter in single-lane roads and yet presents desirable performance in more complex multiple-lane scenarios. A small deterioration in performance has been noticed in highly dense scenarios which is due to inherent characteristics of the 802.11p protocol when broadcasting messages. As future work, we propose applying and evaluating power control mechanisms to further improve the performance of the protocol.

REFERENCES

- [1] E. M. van Eenennaam and G. Heijenk, "Providing over-the-horizon awareness to driver support systems," in *Proceedings of 4th IEEE Workshop on Vehicle to Vehicle Communications (V2VCOM)*, 2008.
- [2] C. J. G. van Driel, "Driver support in congestion : an assessment of user needs and impacts on driver and traffic flow," Ph.D. dissertation, Enschede, November 2007.
- [3] J. Mittag, F. Schmidt-Eisenlohr, M. Killat, J. Härrä, and H. Hartenstein, "Analysis and design of effective and low-overhead transmission power control for vanets," in *VANET '08: Proceedings of the fifth ACM international workshop on Vehicular Ad-Hoc Networking*. New York, NY, USA: ACM, 2008, pp. 39–48.
- [4] S. Dietzel, B. Bako, E. Schoch, and F. Kargl, "A fuzzy logic based approach for structure-free aggregation in vehicular ad-hoc networks," in *Proceedings of the sixth ACM international workshop on Vehicular Ad-Hoc Networking*. ACM, 2009, pp. 79–88.
- [5] N. Shibata, T. Terauchi, T. Kitani, K. Yasumoto, M. Ito, and T. Higashino, "A method for sharing traffic jam information using inter-vehicle communication," *Proceedings of V2VCOM*, 2006.
- [6] S. Ni, Y. Tseng, Y. Chen, and J. Sheu, "The broadcast storm problem in a mobile ad hoc network," in *Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*. ACM New York, NY, USA, 1999, pp. 151–162.
- [7] E. M. van Eenennaam, "Providing over-the-horizon awareness to driver support systems by means of multi-hop ad hoc vehicle-to-vehicle communication," Master's thesis, University of Twente, The Netherlands, December 2008.
- [8] N. Wisitpongphan, O. K. Tonguz, J. S. Parikh, P. Mudalige, F. Bai, and V. Sadekar, "Broadcast storm mitigation techniques in vehicular ad hoc networks," *IEEE Wireless Commun. Mag.*, vol. 14, no. 6, pp. 84–94, December 2007.
- [9] E. M. van Eenennaam, "microslotted 1-persistence flooding in vanets," in *W3 workshop*, Enschede, The Netherlands, November 2009.
- [10] IEEE, "802.11p draft 2.0," *Draft Amendment to Wireless Access in Vehicular Environments*, November 2006.
- [11] J. Blum and A. Eskandarian, "Avoiding timeslot boundary synchronization for multihop message broadcast in vehicular networks," *IEEE 69th VTC2009-Spring, Barcelona, Spain*, April 2009.
- [12] W. Drytkiewicz, S. Sroka, V. Handziski, A. Koepke, and H. Karl, "A mobility framework for omnet++," in *3rd International OMNeT++ Workshop*, Technische Universität Berlin, Germany, 2003.
- [13] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Journal of Physical Review E (Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics)*, vol. 62, no. 2, pp. 1805–1824, 2000.
- [14] L. Quadstone, "Quadstone paramics v5.0 modeller user guide," *Scotland, UK*, 2004.