A Simple and Robust Dissemination Protocol for VANETs

Ramon S. Schwartz*, Rafael R. R. Barbosa†, Nirvana Meratnia*, Geert Heijenk† and Hans Scholten*
*Pervasive Systems (PS) / †Design and Analysis of Communication Systems (DACS)
Faculty of Electrical Engineering, Mathematics and Computer Science (EWI), University of Twente
PO Box 217, 7500 AE Enschede, The Netherlands
{r.desouzaschwartz, r.barbosa, n.meratnia, geert.heijenk, hans.scholten}@utwente.nl

***Abstract*—Several promising applications for Vehicular Ad-hoc Networks (VANETs) exist. For most of these applications, the communication among vehicles is envisioned to be based on the broadcasting of messages. This is due to the inherent highly mobile environment and importance of these messages to vehicles nearby. To deal with broadcast communication, dissemination protocols must be defined in such a way as to (i) prevent the so-called broadcast storm problem in dense networks and (ii) deal with disconnected networks in sparse topologies. In this paper, we present a *Simple and Robust Dissemination (SRD)* protocol that deals with these requirements in both sparse and dense networks. Its novelty lies in its simplicity and robustness. Simplicity is achieved by considering only two states (cluster *tail* and *non-tail*) for a vehicle. Robustness is achieved by assigning message delivery responsibility to multiple vehicles in sparse networks. Our simulation results show that SRD achieves high delivery ratio and low end-to-end delay under diverse traffic conditions.**

Keywords: Ad-Hoc Networks, VANET, Ubiquitous Computing, Dissemination Protocol.

I. INTRODUCTION

Within the context of ubiquitous computing, there is great interest in the area of Vehicular Ad-hoc Networks (VANETs). Although the name may suggest otherwise, VANETs differ in many aspects from Mobile Ad-hoc Networks (MANETs) [1]. The mobility of vehicles is constrained to single or multiple-lane roads. The movement of vehicles usually follows a common pattern, i.e., same or opposite directions. Due to the fact that the dynamics of vehicles often do not permit the proper establishment of end-to-end paths routing solutions proposed for MANETs, such as Dynamic Source Routing (DSR) [2] and Ad Hoc On-Demand Distance Vector Routing (AODV) [3], have been shown to be unsuitable for vehicular networks [4]. These characteristics prompt the need for new and specific solutions in vehicular environments. An example is the effort to define the IEEE 802.11p standard [5], a Carrier Sense Multiple Access (CSMA) MAC protocol specifically adapted for VANETs.

Applications envisioned for VANETs are generally related to safety, transport efficiency, and information/entertainment [6]. To provide the means to serve various applications, the challenges that arise from the unique environment found in roadways first need be addressed [7]. In particular, safety information must be distributed along the road in such a manner that with minimum end-to-end multi-hop delay it reaches a large number of vehicles. Although it is expected that infrastructure will support some applications, others will have to rely on decentralized communication. In this case, broadcast messages are typically sent to all the surrounding vehicles without any coordination. Therefore, the problem dealt with in this paper is the dissemination of safety messages reliably, timely, and efficiently using vehicle-to-vehicle communications. We will do this assuming that safety messages are distributed to upstream vehicles up to a certain distance.

To deal with broadcast communication, various dissemination strategies must be defined according to the current network situation. In dense networks, the number of broadcasts must be minimized to avoid excessive redundancy, contention and collision rates [1]. Taken together, these undesired factors are referred to as the *broadcast storm problem*. This minimization can be achieved by means of *broadcast suppression techniques* [8]. In sparse networks, on the other hand, a *store-carry-forward* communication model (common in *Delay Tolerant Networks (DTN)* [9]) can take advantage of the mobility of nodes to transfer messages when nodes are geographically separated. In this case, the mobility of vehicles in the opposite direction can be used to improve the delivery ratio in disconnected networks [4], [10].

The contribution of this paper is that we combine an optimized broadcast suppression technique with a store-carry-forward model in a single dissemination protocol, which can operate in both dense and sparse networks. Furthermore, we design an improved store-carry-forward mechanism that provides a high delivery ratio, even for partitioned VANETS. Finally, we provide a first evaluation of the proposed Simple and Robust Dissemination (SRD) protocol.

The remainder of this paper is organized as follows. Section II provides an overview and comparison between SRD and other proposals found in the literature with respect to suppression techniques and protocols that also deal with sparse networks. Next, Section III describes the SRD protocol in detail. Section IV describes the performance evaluation of the

---

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

protocol carried out by means of simulations. Finally, Section V concludes this paper and outlines our future plans.

II. RELATED WORK

Various solutions for VANETs have been proposed to cope with message dissemination under different traffic conditions. In dense scenarios, *suppression* techniques have been proposed to address the so-called broadcast storm problem. Unlike MANETs, only a few suppression techniques have been proposed specifically for VANETs. In [8], three broadcast suppression techniques are presented to be employed in the network layer. Among these three techniques, the Slotted 1-Persistance has achieved the best performance in terms of reducing the number of unnecessary broadcasts while still achieving a low end-to-end delay and high delivery ratio. This technique is time-based and non-probabilistic. Given a fixed number of time slots, the most distant vehicles in the message direction from the source vehicle, i.e., from where the message has been originated, will be given the earliest time slot to rebroadcast. Vehicles assigned to other time slots would then have time to cancel their transmissions upon the receipt of an echo. This would be an indication that the information has already been disseminated and any redundant rebroadcast can be *suppressed*.

However, the Slotted 1-Persistance technique suffers from a synchronization problem [11], [12] that can occur when multiple vehicles are assigned to a single time slot and start their transmissions simultaneously. This results in a substantial deterioration with respect to delivery ratio due to a higher number of collisions. In this work, we tackle this problem by proposing the *Optimized Slotted 1-Persistance* technique.

The use of vehicles moving in the opposite direction to help in message dissemination for sparse networks has been previously studied [10], [13]–[16]. In [13], three moving cases are considered: vehicles moving in the same or the opposite direction of the originator of the message and vehicles moving in both directions. Simulation results demonstrate that the use of vehicles moving in opposite directions improves the dissemination performance in many different scenarios. The Directional Propagation protocol [14] allows directional propagation of messages from a given point of origin. It requires the adoption of a cluster creation/maintenance mechanism and differentiates between inter and intra cluster communication. Further evaluation of this protocol in [10] has shown that vehicle mobility can be used to improve message propagation in scenarios in which conventional MANET protocols would fail due to the lack of end-to-end connectivity. In [16] Abiding Geocast disseminates accident or congestion information to every vehicle passing through a warning zone during the event lifetime.

The Distributed Vehicular Broadcast (DV-CAST) protocol [15] and the Acknowledged Parameterless Broadcast in Static to Highly Mobile (ackPBSM) [17] are the closest to the protocol described in our work. The goal of the DV-CAST protocol is to adapt to different traffic densities, e.g., light traffic, moderate traffic, or traffic jam, while introducing a low overhead in high density situations and managing communication gaps in low density situations. Unlike in our approach, DV-CAST relies on the periodic exchange of *hello messages* between all communicating vehicles. Especially in dense and dynamic networks, if not coordinated properly, hello messages might increase collision and contention, thus wasting bandwidth. Although our approach also requires the exchange of periodic messages, employing a suppression technique prevents the so-called broadcast storm problem and reduces the number of broadcasts. Our approach also avoids the dependency on a single vehicle when bridging radio gaps in the network. All vehicles in the range of the vehicle positioned at the tail of a cluster act as *backup vehicles*. Moreover, in [17] the DV-CAST protocol is reported to have a low reliability. This can be partially explained by unforeseen situations such as overtaking while determining the current traffic density. As in our protocol a vehicle simply needs to determine if it is the tail in a message direction or not, it does not suffer from this problem.

The ackPBSM protocol relies on the use of Connected Dominating Sets (CDS) to perform the broadcast. The protocol constructs the CDS based on 1-hop position information, exchanged via periodic hello messages. Acknowledge information is included in these messages in order to increase the delivery ratio and reduce transmission redundancy. The protocol is designed to operate in both highway and urban scenarios. The lack of message direction is a basic difference between ackPBSM and our protocol. In ackPBSM, all vehicles in the network are intended recipients of a broadcast. We argue that in the context of vehicular communication, a *directional* broadcast is more suitable. For instance, consider a highway with two directions. If an accident occurs in one direction, this information is only relevant to vehicles moving in that direction and that have not yet reached that location. For vehicles that have already passed this location or are moving in the other direction, this information is irrelevant. Similarly to DV-CAST, ackPBSM relies on the use of periodic hello messages, which might also lead to increased collisions and contention. Moreover, the collision of hello messages could cause the protocol to perform erratically, as 1-hop position information would not be up-to-date and acknowledgments would be missing. The protocol also needs to maintain a very large data structure consisting of one timer and two lists per broadcast message, and one extra timer per vehicle in one of these lists. In contrast, our protocol only requires that each vehicle maintain a single timer.

III. PROTOCOL DESCRIPTION

Our Simple and Robust Dissemination (SRD) protocol aims to deliver safety information in both dense and sparse networks. More importantly, it aims to achieve a high delivery ratio with a low propagation delay. For this purpose, we make use of the following approaches to establish that the protocol behavior is adequate for both sparse and dense networks:

- In dense networks, messages should be relayed using the minimum number of vehicles by means of *broadcast* *suppression techniques.*

- In sparse networks, the *store-carry-forward* communication model is employed to deliver messages whenever multi-hop connectivity among vehicles is not available.

A. Concept Definitions

To better understand the problem, the following definitions are used during the explanation of our protocol in the following sections.

Definition 3.1 *(Vehicle Cluster). Given a multiple lane road, where vehicles move in both easterly and westerly directions, a vehicle cluster is defined as a group of vehicles with multi-hop radio connectivity at a time instant t.*

Definition 3.2 *(Message Direction). Given a multiple lane road, where vehicles move in both easterly and westerly directions and a source vehicle generating a new message moves in direction d, the message direction is denoted as the direction opposite to d.*

Definition 3.3 *(Cluster Tail). Given a vehicle cluster vc and a message direction d, the cluster tail is defined as the vehicle within cluster vc with no radio connectivity with other vehicles positioned further in message direction d, i.e., the final vehicle belonging to vc in message direction d.*

Definition 3.4 *(Gap). Given two vehicle clusters vc_1 and vc_2, message direction d, and cluster tail ct_1 of cluster vc_1, the gap is defined as:*

$$Gap = D(vc_1, vc_2) - CR(ct_1) \tag{1}$$

where $D(vc_1, vc_2)$ is the relative distance between clusters vc_1 and vc_2, i.e., the distance between ct_1 and the first vehicle in the message direction d of vc_2. $CR(ct_1)$ is denoted as the communication range of vehicle ct_1.

B. Requirements and Assumptions

For a proper operation of the SRD protocol, we require that vehicles are able to determine their position on the road using, for example, the Global Positioning System (GPS). It is not necessary that every vehicle is equipped with wireless communication devices. If they are, however, it is required that the radio ranges provided by these devices are symmetric and working at the same radio frequency, i.e., if a vehicle C_1 can communicate with a vehicle C_2, a transmission from C_2 will also reach C_1.

We also require the existence of at least one application running on top of SRD that generates messages that are periodically aimed at each road direction. Within this paper, this type of application is referred to as *periodic application*. In such an application, a timer for each direction guarantees that a new message is generated and sent upon its expiration or reset upon the receipt of a new message flowing in the target direction. Ideally, the leading vehicles in each direction will always be responsible for generating new messages. SRD uses the existence of periodic messages to determine whether a vehicle is the cluster tail or not. In the context of safety-related applications, two types of messages are assumed to exist: *beacon* and *event-driven* messages [18]. The first type comprises periodic preventive safety messages that are meant to convey information about the state of the source vehicle, i.e., position, speed, etc., and possibly also aggregated data regarding the state of its neighbors. An example of applications generating such beacon messages is presented in [11], where an approach to providing the driver with information about the traffic conditions, e.g., the existence of traffic jams ahead in the road, is described in detail. Event-driven messages are triggered upon the detection of a hazard, e.g., hard braking of cars in front. Critical messages of this type must be delivered as quickly as possible along the road. We argue that such messages should also be retransmitted periodically and maintained flowing for a predefined duration time in the event area in order to increase the chance that at least one copy of each message is received by all vehicles driving towards the critical region [16].

The broadcast messages are required to have an ID that enables vehicles to distinguish different broadcasts and to identify rebroadcasts. It is also necessary that the message also contains the position of the sender and a timestamp. We assume that messages have an expiration mechanism which could be based on distance (to prevent the receipt of messages originated hundreds of kilometers away), and/or time (to prevent the sending of old messages). A message could be considered expired when, for example, it reaches the end of a highway or simply after it reaches vehicles more than 10 km away from the event. Finally, we assume that no roadside infrastructure is available. Although devices alongside the road could help, e.g., identify communication radio gaps, their deployment could be too costly and it would usually take years to build such infrastructure.

C. Optimized Slotted 1-Persistence

Because of the best performance achieved among the techniques proposed in [8], the suppression technique considered in this work is based on the Slotted 1-Persistance; however, with a slightly altered formula to guarantee an equal distribution of vehicles among the time slots assigned[1]. More importantly, we introduce an additional delay within each time slot to cope with the synchronization problem found in [11], [12] to reduce the number of collisions that may occur during a rebroadcast.

The time slot assignment in our technique is called *Optimized Slotted 1-Persistence* and is defined as follows. When vehicle j moving in the 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.

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

---

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

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 *minimum* function is necessary, since the transmission range R is an estimate based on the power level employed and vehicles in reality could be positioned at farther positions.

The time slot number S_{ij} assigned to vehicle j is defined by the following equation:

$$S_{ij} = \lfloor NS \times (1 - PD_{ij}) \rfloor \tag{3}$$

where NS is the total number of time slots utilized. If vehicles are uniformly distributed within the transmission range of vehicle i, they will be equally distributed among the NS time slots reserved. S_{ij} will vary within the interval [0, $NS - 1$].

In most vehicular applications, a message may only be generated in one road direction and may concern vehicles upstream of that particular direction. For instance, when an accident occurs on the road, vehicles upstream must be aware of it and this warning message may not be equally important to vehicles driving in the opposite direction. Therefore, later time slots can be given to them. In this way, by assigning fewer vehicles to each time slot, the number of message collisions during rebroadcasts can be reduced effectively. An improvement we propose when dealing with both directions of a particular road is to double the total number of time slots utilized and to give the second half (later) of time slots to vehicles moving in the opposite direction, as defined by the following condition:

if *vehicle's direction ≠ information origin's direction* **then**
$S_{ij} \leftarrow S_{ij} + \frac{NS}{2}$

In this way, the time slot range is equally divided in [0, $\frac{NS}{2} - 1$] for the higher priority direction and [$\frac{NS}{2}$, $NS - 1$] for the opposite direction.

The time that vehicles have to wait before rebroadcasting at time slot S_{ij} is calculated by equation 4:

$$T_{Sij} = S_{ij} \times st \tag{4}$$

where the slot time st is a value larger than the one-hop delay that includes the medium access delay, transmission delay and propagation delay.

Assigning vehicles to different time slots clearly breaks the synchronization present in the simple flooding approach, where all nodes would rebroadcast simultaneously upon the receipt of a message. The slot time st is defined in such a manner that it gives vehicles assigned to later time slots the opportunity to cancel their transmissions, since the message has already been rebroadcast. Therefore, ideally only vehicles assigned to the earliest time slot would rebroadcast. However, a similar synchronization on a smaller scale can still occur when multiple vehicles are assigned to a single time slot and start their transmission simultaneously. Such a synchronization problem has been identified in [11]. To cope with this problem, a variation of the slotted 1-Persistence technique called microSlotted 1-Persistence Flooding has been proposed. The proposed scheme functions in the same way as the Slotted 1-Persistence Broadcasting scheme but with a small additional delay, i.e., the *micro slots*, within each time slot to break the defined synchronization. The same problem has been identified and referred to as the *Timeslot Boundary Synchronization Problem* in [12]. Differently, such work describes design guidelines for extra measures to be taken not only in the network layer but also in the link layer by inserting a pseudo-random delay to SIFS in the IEEE 802.11p MAC layer. Especially in congested networks, an additional delay introduced uniquely in the network layer does not suffice when nodes experience high contention in the link layer, as their timeslots could be again aligned.

As in [12], we support the position that the synchronization must be broken in both the network and link layers to be completely effective. However, as a preliminary solution we follow the guidelines proposed in [12] but 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 those guidelines, the extra delay must be chosen from a near continuous interval in order to completely break the alignment of timeslot boundaries instead of using fixed-size micro time slots as proposed in [11]. The additional delay AD_{ij} is then defined as follows:

$$AD_{ij} = \begin{cases} D_{max} \times (1 - PD_{ij}) & \textbf{if } vehicle_{dir} = info_{dir}; \\ D_{max} \times (2 - PD_{ij}) & \textbf{if } vehicle_{dir} \neq info_{dir}. \end{cases} \tag{5}$$

where D_{max} is the maximum delay allowed, $vehicle_{dir}$ is the vehicle's direction, and $info_{dir}$ is the information origin's direction. Following the idea adopted for the assignment of time slots, vehicles driving in the same direction where the information has been originated receives smaller delay values than vehicles driving in the opposite direction. The result is that for each road direction 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 expressed in equation 6.

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

D. The Protocol

The SRD protocol decision tree diagram is shown in Figure 1. In the tail state, a vehicle *stores* all broadcasts received and rebroadcasts them with the flag *FromTail* set to true. The tail is responsible for *carrying* these messages until the connectivity in the message direction is established. The tail then *forwards* its stored messages, in this way concluding the *store-carry-forward* mechanism. Vehicles in the non-tail state have two responsibilities: (i) when receiving a broadcast message, if the recipient vehicles are far away in the message direction, the message is simply rebroadcast using the Optimized Slotted 1-Persistence technique. If vehicles are not further in the message direction upon the receipt of a message, previously scheduled messages with the same ID are suppressed and the message is dropped (if not sent by the tail). Here, it is assumed that the message has already been propagated towards the region of interest and does not need to be rebroadcast; (ii) vehicles in this state also store messages sent by the tail (with the *FromTail* flag set to true). This is especially important for improving the protocol reliability as we show later on.

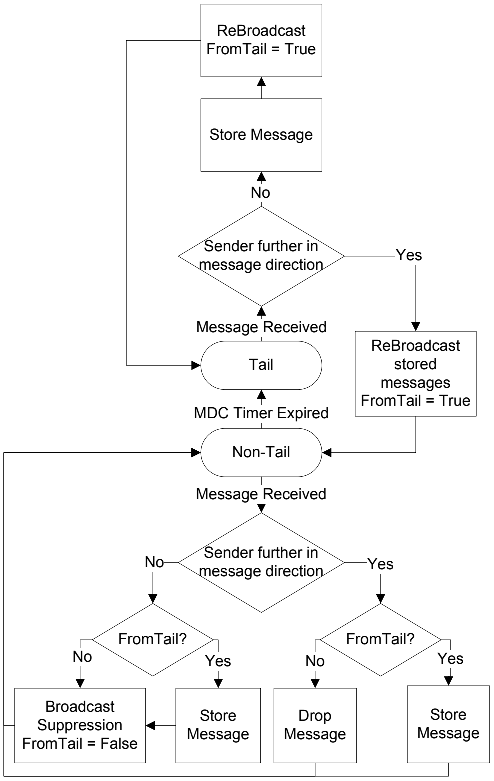

Figure 1. SRD Protocol Decision Tree.

Transitions between the states occur as follows. A vehicle goes from the non-tail to the tail state when it goes for longer than the Message Direction Connectivity (MDC) timer value without receiving a message retransmission from a vehicle further away in the message direction. The MDC timer duration should be defined in such a way that the movement of vehicles is negligible, e.g. tens of milliseconds, and it is dependent on the maximum possible time for a message transmission (i.e. it should take into account message collisions, the *exponential backoff* mechanism, the broadcast suppression technique used, etc.). In the remaining transition, a vehicle moves from the tail to the non-tail state upon the receipt of any message from a vehicle further away in the message direction.

The basic operation of the SRD protocol is shown in Figure 2. In Figure 2(a), a periodic application running in every vehicle aims to generate messages toward the East. When a message coming from the West and disseminated by SRD reaches and is rebroadcast by vehicle S, all vehicles except for the tail simply rebroadcast the message generated using the proposed broadcast suppression technique. When a non-tail vehicle receives a message from another non-tail vehicle that is further in the message direction, it simply drops the message and cancels (suppresses) any previously scheduled transmission in case the message received is an echo. The direction in which vehicles are moving is irrelevant, as all vehicles in the message direction help to disseminate the message. Whenever the broadcast message reaches the tail (C_0 in Figure 2(a)), the cluster tail stores the message and rebroadcasts it with the *FromTail* flag set to true. However, all non-tail vehicles that hear the rebroadcast from the cluster tail would also store the message.

A change in the cluster tail is shown in Figure 2(b), in which C_1 listens to a rebroadcast from the tail C_0. Even though C_1 realizes that the sender was not further in the message direction, the message is stored as it comes with the *FromTail* flag set to *True*. Following the protocol, C_1 rebroadcasts it using the broadcast suppression technique. This rebroadcast is needed since C_1 does not yet know whether it is the new tail. C_0 then receives this retransmission and verifies that the sender is further in the message direction. Consequently, it retransmits all stored messages and performs the transition to the non-tail state. This retransmission is done to cover two possibilities. First, there could be a gap after C_1 further in the message direction and C_1 would become the new cluster tail (as shown in Figure 2(b)). In this case, the rebroadcast is done to guarantee that the new tail has a copy of all messages from the old tail (C_0). In the second case (not shown in the figure), the gap does not exists, i.e. there is a vehicle in the range of the C_1 that is not in the range of C_0. The retransmission in this case will cause C_1 to relay all messages to this further vehicle and consequently to all others that it might be connected to.

As C_1 is moving further in the message direction, at some point it enters in the communication range of C_2, reaching a new cluster, as shown in Figure 2(c). When this happens, C_1 eventually receives a message from C_2. As C_2 is further in message direction, C_1 makes the transition from tail to non-tail state, rebroadcasting every stored message it carries. At this point, C_2 and all non-tail vehicles within its cluster will rebroadcast the messages received in order to spread them to other vehicles further in the message direction.

Every time the tail receives a message, it not only stores the message, but it also retransmits the message with the *FromTail* set to true. By doing so, all vehicles in the range of the tail will also have a copy of that message. If the tail fails or turns off the road, eventually another vehicle will become the new tail. Since such vehicle would already have a copy of all messages received from the old tail, it can rebroadcast them whenever the MDC is restablished. Message delivery thus is not dependent on a single vehicle. In the example shown in

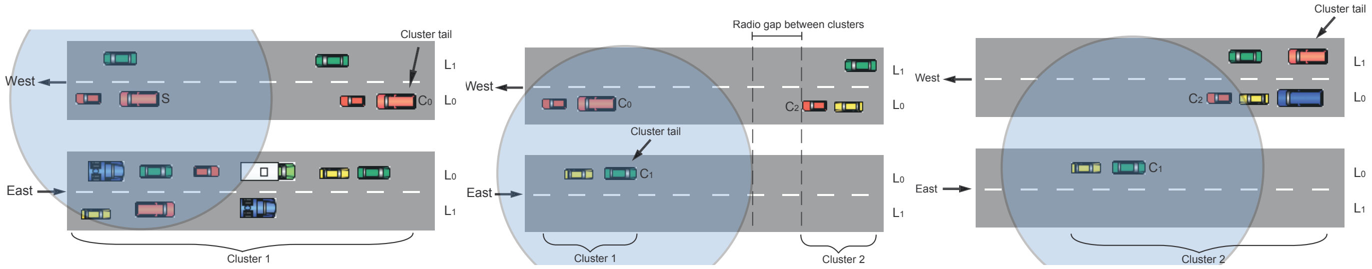

(a) time = t_0 (b) time = $t_0 + x_1$ (c) time = $t_0 + x_1 + x_2$

Figure 2. Protocol Description

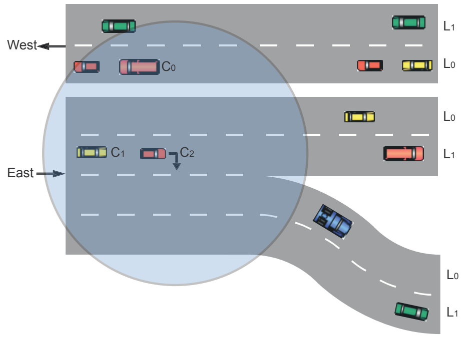

Figure 3. Robustness Motivation

Figure 3, the tail C_2 turns off the road, causing C_0 or C_1 to make the transition to the tail state. As they both have copies of C_2's messages, whichever one makes the transition will be able to retransmit the messages once the MDC is reestablished.

One important remark regarding broadcasting efficiency is that since the rebroadcast from the tail is always required, the tail has a higher priority in the broadcast suppression technique, in order to avoid redundant retransmissions from non-tail vehicles within that region. This priority is implemented by reserving the first timeslot to the tail in the Optimized Slotted 1-Persistence technique and with a smaller additional delay within the time slot when compared with other vehicles.

IV. PERFORMANCE EVALUATION

In this section, we present the performance evaluation of the SRD protocol carried out by means of simulations with Omnet++ 4.0. Our goal is two-fold: (i) verify the advantages of employing the Optimized Slotted 1-Persistence compared to the original Slotted 1-Persistence protocol; and (ii) study the consequences of using both directions of a road with SRD to disseminate information upstream. To achieve these goals, we evaluate the three protocols under both sparse and dense scenarios.

In our simulations, we utilize the Mobility Framework[2] and adjust the available implementation of the IEEE 802.11b protocol to comply with basic specifications of the 802.11p version. In the MAC layer, we set the bit rate to 6 Mbit/s, the Contention Window (CW) to values between 15 and 1023, the slot time to 13 μs, the SIFS to 32 μs, and the DIFS to 58 μs. In the physical layer, we operate on the 5.9 GHz frequency band, with 10 MHz of bandwidth. Based on estimates, we set the transmission power to 168.98 mW to achieve 500 meters of interference range and 250 of transmission range, assuming the Friis Free Space propagation model.

For the suppression technique mechanism, we set st to 5 ms (as proposed in [8]) and define the total number of time slots NS to 5 for the Slotted 1-Persistence protocol and to 10 (5 for each road direction) for the Optimized Slotted 1-Persistence protocol. For the maximum additional delay D_{max}, we use 1 ms. The MDC timer defined in the SRD protocol is set to expire after 60 ms. This accounts for 10 time slots (10 × 5 = 50 ms) plus possible extra delays introduced by an excessive busy medium.

Our evaluation considers the following metrics:

- ***Delivery Ratio***: the percentage of vehicles which receive each message broadcast. Ideally, dissemination protocols must achieve a delivery ratio percentage close to 100%.
- ***Total Channel Utilization***: the percentage of time spent by an arbitrary vehicle transmitting and receiving messages with respect to the total simulation time. In addition to the transmission time, the channel utilization takes into account any noise detected by a vehicle, i.e., errors or collisions during message receptions. This metric evaluates how efficiently the medium is utilized by each protocol.
- ***Delay***: the total time taken for a message to propagate from one end to the other of the road length considered in each scenario. This is particularly important for critical information that must be disseminated as rapidly as possible.

A. Static Scenarios

In the first set of scenarios, we study the performance of the protocols for various traffic densities. To allow that, we simulate a two kilometer road with vehicles placed in both Easterly and Westerly directions, with each direction comprising two lanes. For the sake of simplicity, we consider scenarios where vehicles are equally spaced in such a way that there is no radio gap between them. We vary the number of vehicles from 20 to 100 vehicles/km/lane in 20 vehicles/km/lane steps. Every 2 seconds, a new message of 500 bytes is generated

---

[2] http://mobility-fw.sourceforge.net

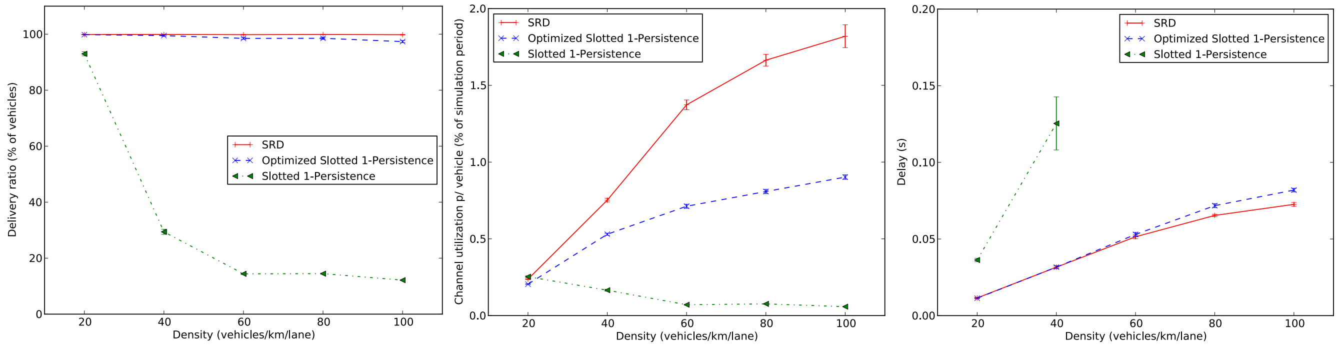

(a) Delivery Ratio x Density (b) Total Channel Utilization x Density (c) Delay x Density

Figure 4. Performance Evaluation for Static Scenarios with 95% confidence intervals

by the foremost vehicles in each direction amounting to 30 messages in a single simulation run of 60 seconds. For each density, we perform 50 runs, accounting for 1500 different messages propagated along each direction. Our goal is to evaluate the scalability of each protocol considered for all metrics described.

In terms of the delivery ratio, Figure 4(a) shows a clear performance deterioration for the Slotted 1-Persistence protocol due to the high number of vehicles assigned to a common time slot and thus a high probability of collision as a result of simultaneous rebroadcasts. The other two protocols, namely Optimized Slotted 1-Persistence and SRD, benefit from the time slot optimization proposed in Section III, which diminishes the probability of collision by including a small additional delay before a rebroadcast. Interestingly, the fact that SRD uses both directions to disseminate information maintains the delivery ratio at almost 100% at all density levels, whereas Optimized Slotted 1-Persistence suffers from a small linear decrease in the delivery ratio from a density of 60 up to 100 vehicles/km/lane. Such improvement can be explained by the higher number of vehicles assigned to the double number of time slots (half for each direction) and participating in each rebroadcast. There is in fact a redundancy in the sense that collisions/errors occurred in rebroadcasts of vehicles assigned to early time slots can be compensated by rebroadcasts from vehicles assigned to later time slots.

However, the higher redundancy – and consequently better delivery ratio – observed in SRD introduces a higher channel utilization compared with the remaining protocols as shown in Figure 4(b). At high densities, the channel utilization observed in SRD is almost double that achieved with Optimized Slotted 1-Persistence. In fact, this could be expected by double the number of vehicles participating in dissemination when both directions are used. Unlike for the remaining protocols, the channel utilization with Slotted 1-Persistence decreases with higher densities. This is simply due to the decrease in delivery ratio observed previously and the consequent lower number of vehicles participating in each rebroadcast. Apart from the divergence between SRD and Optimized Slotted 1-Persistence, we can note that a message generation periodicity of 2 seconds occupies less than 2% of the total channel utilization time, thereby creating the opportunity for other applications to run simultaneously with SRD.

Finally, we study the average end-to-end delay in both road directions. Figure 4(c) shows that both SRD and Optimized Slotted 1-Persistence protocols achieve similar performances for different densities, with a sightly better performance observed with SRD at higher densities due to the redundancy introduced and more vehicles attempting to rebroadcast. We can observe that the delay is generally increased at higher densities for all protocols. Under scenarios where there are vehicles assigned to every time slot, the high probability of collisions and errors is the main cause for such an increase. Because the messages disseminated with the Slotted 1-Persistence protocol could only complete their end-to-end path at low densities, only results for densities of 20 and 40 vehicles/km/lane are shown in the figure. In addition to the low delivery ratio found, the synchronization problem explained in the Slotted 1-Persistence protocol also results in a higher end-to-end delay due to an even higher number of collisions compared with that observed from the other two protocols.

Overall, the results obtained indicate Optimized Slotted 1-Persistence outperforms the original Slotted 1-Persistence in every metric considered. Furthermore, the use of both road directions with SRD improves the delivery ratio and reduces delay but at the cost of a higher channel utilization.

B. Mobility Scenarios

In the second set of scenarios, we focus on evaluating the effects of using both road directions to disseminate messages under mobility scenarios. Due to the poor results achieved by the Slotted 1-Persistence protocol, we concentrate on the comparison of SRD and Optimized Slotted 1-Persistence protocols in various scenarios that include radio gaps between vehicle clusters. Our goal is to study the performance in terms of delivery ratio and channel utilization of both protocols and validate the advantages of using SRD in scenarios with vehicles at high speed and with intermittent multi-hop path disconnections. To achieve this goal, we consider four scenarios in a highway that has two lanes per road direction and

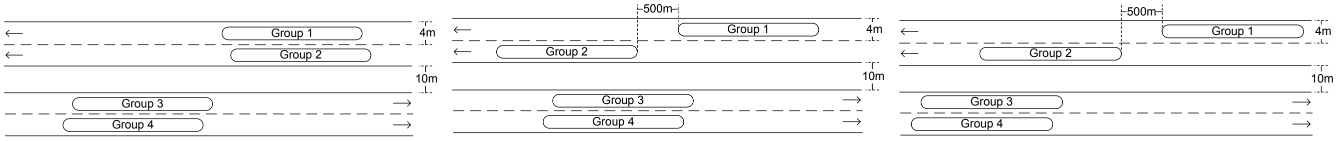

(a) Scenarios 1 and 2 (b) Scenario 3 (c) Scenario 4

Figure 5. Mobility Scenarios

four vehicle clusters (groups). Lanes are 4 meters wide with a 10 meter space between the directions. Similarly to previous simulations under static scenarios, a new message of 500 bytes is generated by the foremost vehicles in each direction at every 2 seconds for each simulation run of 60 seconds.

In Scenarios 1 and 2, all four lanes are very busy, with 100 vehicles/km/lane. In each lane there is a group of 250 vehicles separated by 10 meters. The initial state is shown in Figure 5(a). The only difference between these two scenarios is the speed at which the vehicles move. In scenario 1, vehicles move at speeds between 2 and 2.5 km/h, while in scenario 2, vehicle speeds vary between 115 and 120 km/h. There is always connectivity between the groups during the simulation time. Scenarios 3 and 4 simulate situations with radio gaps between vehicles clusters. In scenario 3, shown in Figure 5(b), there is a 500 m gap between groups 1 and 2 in such a way that they cannot communicate directly. SRD transparently uses the opposite direction to relay messages from group 1 to group 2, thereby bridging the gap. In this scenario, each group has a density of 20 and vehicles move at speeds between 115 and 120 km/h. In scenario 4 (Figure 5(c)) the gap also exists but there are no vehicles moving in the opposite direction in the initial state. To bridge the gap, vehicles moving in the opposite direction use the store-carry-forward mechanism. Vehicle densities and speeds are the same as in scenario 3.

Figure 6 illustrates the delivery ratio obtained in each scenario using both protocols. The results obtained reflect only the delivery ratio achieved for vehicle groups 1 and 2 in order to evaluate the advantages of using both directions when radio gaps are present. In scenarios 1 and 2, both protocols achieve delivery ratios of between 95% and 98%. Notably, the high speed of vehicles in Scenario 2 does not affect the performance of any of the protocols. In fact, as shown in Figure 4(c), during the period a message propagates from one end to the other, i.e., a maximum of 0.07 seconds, vehicles practically do not move: at 120 km/h, vehicles would only move 2.34 meters. As observed for static scenarios and for the same reasons previously explained, the use of both road directions in the SRD protocol results in a slight improvement over Optimized Slotted 1-Persistence. In scenarios 3 and 4, we can verify that SRD outperforms the Optimized Slotted 1-Persistence protocol through utilization of both directions directly in scenario 3 and using the store-carry-forward approach in scenario 4. Since the size of vehicle groups 1 and 2 are the same, Optimized Slotted 1-Persistence can only achieve at most 50% of delivery ratio in both scenarios. The use of the store-carry-forward approach in scenario 4 results in slight decrease of the delivery ratio (90%) in SRD. This decrease in delivery ratio can be explained by the fact that in SDR cluster tails do not wait for confirmation (an echo message) from the other vehicles at the time when two clusters meet. This could be easily overcome with a simple message echo verification and the resending of messages in case of failures.

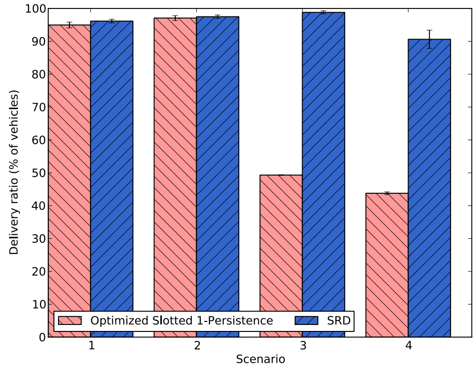

Figure 6. Delivery Ratio x Mobility Scenarios with 95% confidence intervals

Figure 7 shows the channel utilization results for each scenario considered. Similarly to the results obtained for static scenarios, the channel utilization is slightly higher with SRD in comparison with Optimized Slotted 1-Persistence in Scenarios 1, 3, and 4. This is also explained by the fact that with SRD vehicles from both directions are used for the message dissemination. In contrast, there is a substantial increase in channel utilization with SRD in Scenario 2. In this scenario specifically, the high density of 100 vehicles/km/lane and the high speed of vehicles result in an close approximation of vehicle groups from both directions. Therefore, the higher number of vehicles interfering with each other leads to a higher channel utilization when compared with the results obtained in the remaining scenarios.

Overall, SRD is effective at improving the delivery ratio in each scenario considered and even more notably in disconnected networks. As observed for static scenarios, the channel utilization is generally increased with SRD due to a higher number of vehicles participating in the message dissemination.

V. CONCLUSION AND FUTURE WORK

In this paper we have presented a dissemination protocol suitable for both sparse and dense VANETs. The use of *suppression* techniques has been motivated and employed in dense

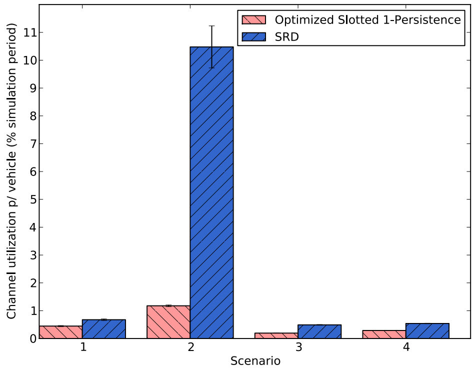

Figure 7. Channel Utilization x Mobility Scenarios with 95% confidence intervals

networks while the *store-carry-forward* communication model has been used in sparse networks. The designed protocol is both *simple* and *robust*. We have proposed an optimization for the suppression technique employed by the SRD protocol. Both optimization and the SRD itself have been evaluated in terms of delivery ratio, channel utilization, and delay. Our experimental results show that SRD can benefit from the use of both directions to disseminate messages along the road. The use of both road directions have increased the delivery ratio considerably at the cost of some extra redundancy and the consequent higher channel utilization. In future work, we will aim to study the effects of such extra redundancy under the presence of other protocols and specially under scenarios with a more realistic vehicle distribution. Furthermore, we intend to compare SRD directly with other approaches such as DV-CAST and ackPBSM. Finally, we will propose measures also in the link layer with regard to the synchronization problem in the time-based suppression technique and compare it with other approaches, such as microSlotted 1-Persistence Flooding.

ACKNOWLEDGMENTS

The authors express their sincere gratitude to Martijn van Eenennaam, Georgios Karagiannis and Wouter Klein Wolterink for their collaboration and for the discussion of developing optimizations for broadcast suppression techniques.

REFERENCES

[1] 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.

[2] D. Johnson and D. Maltz, "Dynamic source routing in ad hoc wireless networks," *Kluwer International Series in Engineering and Computer Science*, 1996.

[3] C. Perkins and E. Royer, "Ad-hoc on-demand distance vector routing," in *Second IEEE Workshop on Mobile Computing Systems and Applications, 1999. Proceedings. WMCSA'99*, 1999, pp. 90–100.

[4] N. Wisitpongphan, F. Bai, P. Mudalige, and O. Tonguz, "On the routing problem in disconnected vehicular ad hoc networks," in *Proc. IEEE Infocom*, 2007.

[5] D. Jiang and L. Delgrossi, "Ieee 802.11p: Towards an international standard for wireless access in vehicular environments," in *Proc. IEEE Vehicular Technology Conference*, 11–14 May 2008, pp. 2036–2040.

[6] H. Hartenstein and K. P. Laberteaux, "Topics in ad hoc and sensor networks - a tutorial survey on vehicular ad hoc networks," *IEEE Communications Magazine*, vol. 46, no. 6, pp. 164–171, June 2008.

[7] M. Torrent-Moreno, M. Killat, and H. Hartenstein, "The challenges of robust inter-vehicle communications," in *Proc. VTC-2005-Fall Vehicular Technology Conference 2005 IEEE 62nd*, vol. 1, 2005, pp. 319–323.

[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 Communications*, vol. 14, no. 6, pp. 84–94, December 2007.

[9] L. Franck and F. Gil-Castineira, "Using Delay Tolerant Networks for Car2Car communications," in *IEEE International Symposium on Industrial Electronics, 2007. ISIE 2007*, 2007, pp. 2573–2578.

[10] A. Agarwal, D. Starobinski, and T. Little, "Exploiting downstream mobility to achieve fast upstream message propagation in vehicular ad hoc networks," *Mobile Networking for Vehicular Environments*, 2007.

[11] 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.

[12] J. Blum and A. Eskandarian, "Avoiding timeslot boundary synchronization for multihop message broadcast in vehicular networks," *2009 IEEE 69th Vehicular Technology Conference: VTC2009-Spring 26–29 April 2009, Hilton Diagonal Mar Hotel, Barcelona, Spain*, 2009.

[13] T. Nadeem, P. Shankar, and L. Iftode, "A comparative study of data dissemination models for vanets," in *3rd Annual International Conference on Mobile and Ubiquitous Systems (MOBIQUITOUS)*, 2006.

[14] T. Little and A. Agarwal, "An information propagation scheme for VANETs," *IEEE Intelligent Transportation Systems, 2005. Proceedings*, pp. 155–160, 2005.

[15] O. Tonguz, N. Wisitpongphan, F. Bai, P. Mudalige, and V. Sadekar, "Broadcasting in VANET," *2007 Mobile Networking for Vehicular Environments*, pp. 7–12, 2007.

[16] Q. Yu and G. Heijenk, "Abiding geocast for warning message dissemination in vehicular ad hoc networks," in *Proc. IEEE Vehicular Networks and Applications Workshop*, 2008.

[17] F. Ros, P. Ruiz, and I. Stojmenovic, "Reliable and efficient broadcasting in vehicular ad hoc networks," in *Vehicular Technology Conference, 2009. VTC Spring 2009. IEEE 69th*, April 2009, pp. 1–5.

[18] J. Mittag, F. Schmidt-Eisenlohr, M. Killat, J. Härri, 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 Inter-NETworking*. ACM, 2008, pp. 39–48.