Towards Lightweight Information Dissemination in Inter-Vehicular Networks

Davide Sormani, Gabriele Turconi,
Paolo Costa, Davide Frey, Matteo Migliavacca and Luca Mottola
Dipartimento di Elettronica e Informazione
Politecnico di Milano, Italy
[costa, frey, migliava, mottola]@elet.polimi.it

ABSTRACT

Vehicular ad hoc networks have recently been proposed as an effective tool for improving both road safety and the comfort experienced while driving. Vehicles may propagate information about potentially dangerous events such as lane changes or sudden slowdowns to vehicles in their vicinity. Moreover they can inform vehicles approaching from farther areas about accidents and possible traffic jams. In both cases, data must be routed to specific areas, along paths determined by the underlying road traffic conditions.

In this paper we propose a novel approach to address this routing problem. First, we define a message propagation function that encodes information about both target areas and preferred routes. Second, we show how this function can be exploited in several routing protocols; and finally, we evaluate the effectiveness of our approach by means of simulation. Results highlight the good performance of our routing approach in sparse as well as in dense networks.

Categories and Subject Descriptors: C.2.1. [Network Architecture and Design]: Store and forward networks, Wireless Communication; C.2.2 [Network Protocols]: Routing Protocols

General Terms: Algorithms, Experimentation.

Keywords: Vehicular Ad Hoc Networks.

1. INTRODUCTION

Anyone who has been stuck in a traffic jam has dreamed about the ability to be notified on traffic congestion in time to decide for a different route. However, existing Traffic Information Systems (TISs) most often fail to provide accurate and timely information to drivers. Infrastructure-based TISs are inherently associated with very high deployment and maintenance costs. This forces road-management companies to make significant investments to render such TISs useful on a large-scale.

For this reason, the availability of low-cost wireless communication devices has fostered increasing interest in infrastructure-less communication among vehicles. The ability to communicate without a fixed infrastructure is in fact likely to facilitate the introduction of TISs by cutting down the fixed costs associated with their deployment. The active research in the field and the involvement of car manufacturing companies suggest that in a very near future wireless communication among vehicles will provide improvements both to road safety and to the comfort experienced while driving. To address safety, vehicles may propagate information about their speed and movement direction to vehicles in their vicinity and inform them about potentially dangerous events such as lane changes or sudden slowdowns. To address comfort, they can propagate information to farther areas and inform approaching vehicles about situations of traffic congestion. Communication among cars may also allow drivers to communicate effectively with fixed stations along their intended routes, for example to query about the availability of parking places in a given area. In all these cases, routing protocols should ideally deliver messages only to specific areas in the road network. For example, information about an accident or a traffic congestion should be propagated only along the streets that lead to the congested area.

In this paper we address the above requirements with a novel approach for disseminating information to a set of target areas while taking into account the structure of the underlying road network. The approach is based on a *propagation function* that encodes the destination areas of each message as well as the roads it should follow to reach such areas.

The paper is structured as follows. Section 2 describes the application scenario we target in our work. Section 3 presents our system model and defines the concept of *propagation function* used throughout the paper. Section 4 describes a set of protocols that exploit a propagation function to route messages in vehicular networks. Section 5 validates the approach by means of simulation. Section 6 places our work in the context of related efforts, and finally, Section 7 concludes the paper.

2. SCENARIO

Let us consider an accident happening on the highway depicted in Figure 1. Two cars crash while travelling southbound in the point indicated by the circle. We envision a system in which vehicles cooperate to inform each other and emergency personnel on the accident. This system prevents dangerous situations for the drivers that are already close to the accident. Moreover it promptly informs drivers that

VANET'06, September 29, 2006, Los Angeles, California, USA.

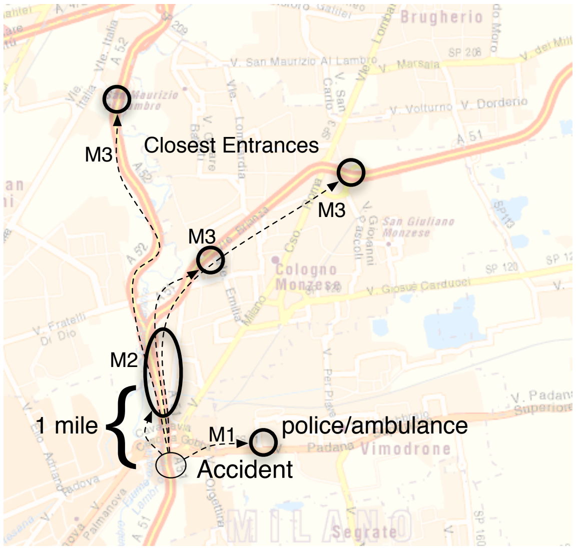

Figure 1: A sample accident scenario on a highway, with messages being delivered to emergency personnel, approaching vehicles, and closest highway entrances.

are farther away and allows them to take alternative routes to reach their destinations.

As soon as the accident happens, the crashed cars and possibly other cars in the vicinity immediately generate a message $M1$ that is routed towards the nearest ambulance and police stations. This guarantees prompt assistance to the occupants of the vehicles involved in the accident. A second message $M2$ is delivered to the vehicles within one mile approaching the accident area and tells them to slow down. Finally a third message $M3$ is propagated to the closest highway entrances north of the accident to inform vehicles of a likely traffic congestion. This allows drivers receiving $M3$ to plan for different routes to their destinations.

A common feature in the propagation of messages $M1$, $M2$, and $M3$ is that they all must be routed to specific locations or areas in the road network. In principle, $M1$ could also be delivered in a unicast fashion to the police and to emergency personnel. However, $M2$ and $M3$ do not have specific destinations. Rather they must reach as many vehicles as possible within their target zones: the preceding mile in the highway for $M2$, and the nearest preceding entrances for $M3$.

A further aspect to consider is that the propagation of these messages should ideally occur along the highway or along other roads in the vicinity. In particular, while vehicles will generally reach their destinations most easily by taking fairly empty streets, message propagation is most effective following streets where density of vehicles is sufficiently high.

In this paper, we start from these observations and describe a novel approach to routing messages towards a set of target zones while taking into account the traffic conditions of the underlying road network. At the core of the approach is a message *propagation function* that allows nodes to route messages through zones characterized by a sufficiently high density of vehicles, thereby increasing their chances of reaching their intended destinations.

3. SYSTEM MODEL

Based on the scenario outlined in Section 2, we are now ready to analyze the characteristics of our approach to information dissemination in Inter-Vehicular networks. In doing this, we consider a network of vehicles communicating over a wireless medium. The communication range of each wireless device may be different for each host. Additionally, we assume that each vehicle has access to some location service such as GPS. The location service provides each vehicle with information about its current position p in a two- or three-coordinate space, that is $p \in C$, with $C \subseteq \Re^2$ or $C \subseteq \Re^3$.

Our routing approach exploits this location service to *deliver messages to vehicles located in a set of specified target zones, following routes determined by road-traffic conditions.* This is achieved by the message originator using a propagation function f defined over the same coordinate space associated with the location service, and by a threshold value v_{th}.

$$f : C \rightarrow \Re$$
$$v_{th} \in \Re$$

The function encodes both the target zones, A^i_{tg}, of a message and the routes that it should follow to reach them. Specifically, the union of all the target zones associated with a given function, $A_{tg} = \bigcup A^i_{tg}$, comprises all the positions in which the value of the function is less than or equal to v_{th}.

$$A_{tg} = \{p \in C | f(p) \leq v_{th}\}$$

The route that a message should take to reach each of the target zones is instead driven by the directions of maximum decrease of the function. More precisely, messages should be steered towards the areas within communication range in which the function returns the lowest values. This associates each target area with an infinity of trajectories that depend on the message's point of origin. Moreover, each propagation function may be associated with several target zones.

Figure 2 shows a propagation function associated with a target zone that is reached by a single major road. The function drives messages along the main road — the black line below the function — and towards the target zone — the black ellipse. It is important to observe that the message originator does not compute a predefined trajectory using the propagation function before sending the message. Rather, the route to the destination is the result of the evaluation of the function at each routing hop. For example, a message that is routed outside the black line[1] in Figure 2 does not need to be routed back towards the line, but it can continue its route along a new trajectory, which still ends up on the desired target area.

The idea behind this routing mechanism is that each target zone acts as a mass that determines a gravitational field. The values of the propagation function can be viewed as the values of the potential associated with this field. Messages should be attracted by the field in the right direction towards decreasing values of the propagation function, that is towards areas of minimum potential.

As we mentioned above, the propagation function is associated with each message by its originator. At an abstract

---

[1] For example because a vehicle responsible for forwarding it has left the main road.

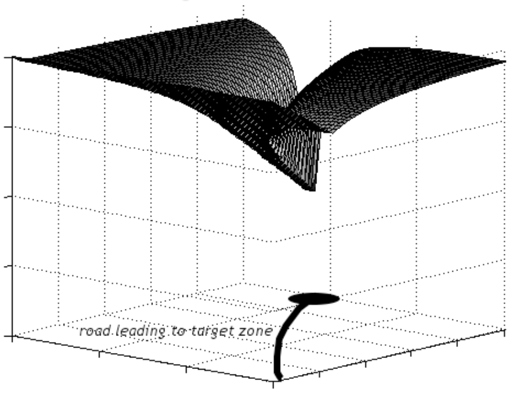

Figure 2: A propagation function and the corresponding target zone.

level, the originator should compute the function based on available information about the structure and traffic conditions of the road network. Nevertheless, the process can easily be optimized by building a library of precomputed propagation functions associated with well-known destination areas. Each function can be determined based on information about average road traffic in different time frames. The resulting library can then easily be stored in each car's navigation system. The originator of a message can thus select the best function from its library depending on the target areas and the current traffic conditions.

The use of a library of functions also makes our approach open to several optimizations. For example nodes along a message's route can replace the function selected by the originator if they have a new version available that reflects more current information. Nevertheless, a detailed analysis of this and other optimizations, is outside the scope of this paper.

A final characteristic of our approach is that message sources can use each well-known function individually or combine and simultaneously use different functions to serve several target areas. In the latter case, the resulting propagation function may exhibit local minima. Therefore routing protocols should take action to minimize the possibility that messages stuck at one minimum become unable to reach the other ones. For example, the application can use propagation functions that contain paths from each source to each destination along which the value of the function is strictly decreasing, or routing protocols can properly employ store-and-forward routing strategies, as described in Section 4. A detailed evaluation of the effectiveness of these techniques in the presence of multiple target areas is the subject of ongoing work and is therefore outside the scope of this paper.

4. PROTOCOLS

Based on the system model and assumptions previously described, we wanted to assess the benefits brought by the propagation function when coupled with different dissemination techniques. In this respect, we do not propose a single, general-purpose routing protocol. Rather, we outline several protocols to explore the impact of the information encoded in the propagation function in various settings, and to investigate the modifications needed in different dissemination techniques. To this end, we proceed from two simple solutions—which will be considered as a baseline for performance comparison—to more sophisticated algorithms featuring probabilistic schemes and store & forward mechanisms. In particular, in probabilistic approaches, the message forwarding decisions are not just deterministic, but instead driven by probabilistic choices. This approach is known to be well suited to the Ad-hoc environment, thanks to its low overhead and good scalability properties [9]. On the other hand, store & forward refers to the ability for a node to locally store a message for a given time period, and later retransmit this message in a location different from where it was originally received. In other terms, this mechanism exploits the hosts in the system as "mules" that physically carry messages [10], hence taking advantage of the nodes' physical mobility. As we will illustrate in Section 5, this enables good performance in terms of message delivery even in sparse networks where connectivity can easily become an issue.

In all the proposed protocols the message transmissions are always broadcast, and forwarding decisions are always taken on the receiver's side. This approach blends well with the highly dynamic scenarios of Inter-Vehicular networks, in that it does not require a proactive maintenance of neighbors' information. Indeed, these may easily become obsolete because of high node speeds. For all protocols it is assumed that, once the target zone is reached, the mechanism turns to flooding propagation as long as the message remains on the target zone.

Common to all proposed approaches is the message format, shown in Figure 3. Most of the fields are self-explanatory. It is just worth noting the difference between the message *Source* and its physical *Sender*. The former is the node that first generated the message, whereas the latter can be any node propagating that message on behalf of the source. Additionally, we approximate the physical space covered by a node's wireless interface as a circular area. In describing the different protocols we will refer to a particular message field as $m.fieldName$, where m is the message and $fieldName$ is one of those in Figure 3. Moreover, we will make use of the functions and procedures defined in Table 1. Finally, we will assume a variable $localPosition$ representing the geographical location of the node processing the message, and each message implicitly dropped if not propagated.

4.1 Baseline Protocols

One Zero Flooding (OZF). Flooding is known to be a basic building block for a wide range of protocols in Ad-hoc networking. To take into account the propagation function, the most simple modification to the basic flooding mechanism is to forward a message only if it is received for the first time *and* the receiver is in a position where the propagation function returns values lower than at the sender's position. This way, one can implement a form of directional flooding where messages are propagated only towards areas where the propagation function returns the lowest values. The core mechanism is illustrated in Pseudocode 1. Notice the forwarding decision is still fully deterministic (hence the name "OneZeroFlooding") and based only on the instantaneous positions of the sender and receiver nodes.

Source	**Timestamp**	**PropagationFunction**	**Sender**	**SenderPosition**	**SenderCommRadius**	**Payload**

Figure 3: The message format common to all protocols illustrated.

Name	**Description**
neverSeen(Message m)	Returns $true$ if m has been received for the fist time, $false$ otherwise.
sendBroadcast(Message m)	Sends a broadcast message to all nodes within the communication range.
randomChoice(Probability p)	Extracts a random number between 0 and 1, and returns $true$ if it is greater than or equal to p, $false$ otherwise.
scheduleTimeoutFor(Message m, ...)	Schedules a timeout event for a given message m. On the expiration of this timeout, the $TimeoutExpired$ procedure will be called with message m as parameter, plus any additional parameter given at the time of scheduling the timeout.
cancelTimeout(Message m)	Cancels a previously scheduled timeout for message m.
evalMovementDirection(Position p)	Evaluates the direction of movement according to the given position and returns a vector representing this direction.
evalGradient(Function f, Position p)	Evaluates the gradient of the given function in the given position, and returns a vector representing this gradient.

Table 1: Functions and procedures used to describe the different protocols.

Pseudocode 1 *One Zero Flooding (OZF)*: on receiving message m.

```
$Function\ f \leftarrow m.PropagationFunction$
if $neverSeen(m)$
$\wedge f(localPosition) < f(m.SenderPosition)$ then
    $m.SenderPosition \leftarrow localPosition$
    $sendBroadcast(m)$
end if
```

Distance-Driven Probabilistic Diffusion (DDPD). Adding a probabilistic choice to OZF requires mapping the benefit associated to forwarding a message onto a [0,1] interval representing the actual forwarding probability. In this case, we represent such a benefit as the geographical distance between the sender and the receiver. Normalizing the distance with respect to the communication radius yields the desired probability. The protocol is detailed in Pseudocode 2. Basically, it is a simple extension to OZF with the addition of a probabilistic decision in place of the fully deterministic scheme.

The two protocols described above represent two baseline solutions, suited for performance comparison. Nonetheless, smarter mechanisms can be devised. These are illustrated next.

4.2 Enhanced Protocols

Function-Driven Probabilistic Diffusion (FDPD). The two previous protocols take into account the propagation function only to recognize the cases where a message is propagating "in the wrong direction", i.e., it is not moving closer to the target zone. However, we can also rely on the values returned by the propagation function to implement a more informed probabilistic scheme, in which higher values of the propagation function correspond to higher chances of forwarding a message.

Pseudocode 2 *Distance Driven Probabilistic Diffusion (DDPD)*: on receiving message m.

```
$Function\ f \leftarrow m.PropagationFunction$
if $neverSeen(m)$
$\wedge f(localPosition) < f(m.SenderPosition)$ then
    $Probability\ p \leftarrow \frac{distance(m.SenderPosition, localPosition)}{m.SenderCommRadius}$
    if $randomChoice(p)$ then
        $m.SenderPosition \leftarrow localPosition$
        $sendBroadcast(m)$
    end if
end if
```

Pseudocode 3 *Function Driven Probabilistic Diffusion (FDPD)*: on receiving message m.

```
$Function\ f \leftarrow m.PropagationFunction$
if $neverSeen(m)$
$\wedge f(localPosition) < f(m.SenderPosition)$ then
    $BestPoint \leftarrow$
        $best(f, m.SenderPosition, m.SenderCommRadius)$
    $Probability\ p \leftarrow$
        $\frac{f(m.SenderPosition) - f(localPosition)}{f(m.SenderPosition) - f(BestPoint)}$
    if $randomChoice(p)$ then
        $m.SenderPosition \leftarrow localPosition$
        $sendBroadcast(m)$
    end if
end if
```

As in the case of DDPD, we need to map the gain associated to forwarding a message onto a [0,1] interval. Based on the above reasoning, we represent this gain with the difference between the evaluation of the propagation function at the sender's and receiver's position. Additionally, to normalize this quantity, we introduce the notion of *best point*. Intuitively, the *best point* is the physical location within the communication radius of the sender node where the propagation function returns the lowest value. In mathematical terms, it is

$$best(f, p, r) = min(f(X)), X \in D(p, r) \qquad (1)$$

where $D(p, r)$ is the physical space covered by the communication radius r of the sender node when in position p. Given the above definition, one can normalize the aforementioned difference with the difference between the evaluation of the propagation function at the *best point* and the same evaluation at the sender's position. The core mechanism of FDPD is then described in Pseudocode 3

Feedback-augmented Store & Forward Diffusion (FSFD). The aforementioned protocols do not make use of store & forward techniques. However, these can be useful in sparse networks where instantaneous connectivity between the sender and the target zone is not guaranteed. In addition, store & forward techniques can be of help in avoiding local minima, dealing with non-convex propagation functions, or avoiding physical obstacles hampering communication, e.g., buildings. In this protocol we propose a simple store & forward technique based on a timeout associated to each received message. Each received message is locally cached, and, when the corresponding timeout expires, the message is re-propagated and the timeout rescheduled. Based on node speed and the actual timeout value, the successive propagations will likely take place in different locations from where the message was originally received.

To stop the periodic triggering of further message propagations, we either reach a maximum number of retransmissions (determined by a protocol parameter), or passively listen for other broadcasts of the same packet by other neighboring nodes. In particular, a timeout is cancelled when the corresponding message is heard from another node lying in a position where the propagation function returns a value lower than at the local position. This mechanism implements a sort of *passive feedback* from other neighboring nodes. In other words, the periodic message propagation is stopped as soon as a node recognizes the presence of another node doing the same from a better position (i.e., lower values) with respect to the propagation function. The core of the protocol is described in Pseudocode 4 (counting the number of retransmissions is not shown for simplicity).

Pseudocode 4 *Feedback-augmented Store & Forward Diffusion (FSFD)*: on receiving message m.

```
$Function\ f \leftarrow m.PropagationFunction$
if $neverSeen(m)$
$\wedge f(localPosition) < f(m.SenderPosition)$ then
    $m.SenderPosition \leftarrow localPosition$
    $sendBroadcast(m)$
    $scheduleTimeoutFor(m)$
else if $\neg neverSeen(m)$
$\wedge f(m.SenderPosition) < f(localPosition)$ then
    $cancelTimeoutFor(m)$
end if
procedure TIMEOUTEXPIRED($Message\ m$)
    $sendBroadcast(m)$
    $scheduleTimeoutFor(m)$
end procedure
```

Function Driven Feedback-augmented Store & Forward Diffusion (FD-FSFD). The previous protocol can be easily extended with the same probabilistic scheme used in FDPD. Basically, instead of blindly rebroadcasting a message, we also perform the same probabilistic choice we described in the case of FDPD. This is based on the notion of *best point*, already described for FDPD. The pseudocode for the protocol is a combination of those in Pseudocode 3 and 4 and is omitted for brevity.

Direction-aware Function Driven Feedback-augmented Store & Forward Diffusion (DFD-FSFD). Store & forward techniques can also take into account the direction of movement so that only nodes moving towards lower values of the propagation function are used to carry messages. To recognize this situation, we evaluate the angle between the direction of node movement and the gradient of the propagation function at the receiving node's position. Specifically, we use the store & forward technique described above only if the absolute value of this angle is less than $\frac{\pi}{2}$. The mechanism is described in Pseudocode 5.

Pseudocode 5 *Direction-aware Function Driven Feedback-augmented Store & Forward Diffusion (DFD-FSFD)*: on receiving message m.

```
$Function\ f \leftarrow m.PropagationFunction$
if $neverSeen(m)$
$\wedge f(localPosition) < f(m.SenderPosition)$ then
    $BestPoint \leftarrow$
        $best(f, m.SenderPosition, m.SenderCommRadius)$
    $Probability\ p \leftarrow$
        $\frac{f(m.SenderPosition) - f(localPosition)}{f(m.SenderPosition) - f(BestPoint)}$
    if $randomChoice(p)$ then
        $m.SenderPosition \leftarrow localPosition$
        $sendBroadcast(m)$
    end if
    $Vector\ v \leftarrow evalMovementDirection(localPosition)$
    $Vector\ g \leftarrow evalGradient(f, localPosition)$
    if $|angle(v, g)| < \frac{\pi}{2}$ then
        $scheduleTimeoutFor(m, p)$
    end if
else if $\neg neverSeen(m)$
$\wedge f(m.SenderPosition) < f(localPosition)$ then
    $cancelTimeoutFor(m)$
end if
procedure TIMEOUTEXPIRED($Message\ m$,
                         $Probability\ p$)
    if $randomChoice(p)$ then
        $sendBroadcast(m)$
    end if
    $scheduleTimeoutFor(m)$
end procedure
```

5. EVALUATION

In this section we evaluate our approach through simulation. The main metrics we adopt to compare the protocols described in Section 4 are *message delivery* and *network traffic*. The former is defined as the ratio between the number of nodes that receive a given message and the overall number of nodes. More precisely, to assess the selectivity of our protocols, we measure the delivery rate both inside and outside the target zone (respectively *delivery-IN* and *delivery-OUT*). Ideally, *delivery-IN* should be kept close to 100%, whereas *delivery-OUT* should be as low as possible. Nonetheless, even in the optimum case, the value of *delivery-OUT* cannot be zero as some nodes outside a message's target zone are required to receive and propagate the message. Network traffic is, instead, measured as the total number of message transmissions, including forwarding operations, across the entire system.

Further insight on the ability of our protocols to steer messages is offered by *diffusion charts*. To provide a snapshot of message propagation, we divide the square area in a grid of 100 cells and fill each cell with 9 additional probing nodes. These nodes are used only to record the fraction of messages received, i.e., they are fixed and cannot forward messages. We send 100 messages from a fixed position in the bottom-left corner, directed to a target zone located in the top-right corner and color each node along the path according to the ratio between the number of messages received by that node and the number of messages generated.

Simulation Settings. All of our simulations are developed with J-Sim [1], a Java-based open-source simulator, integrated with the *Manhattan Mobility Model* [2]. This model enables a more realistic evaluation because it emulates the movement pattern of mobile nodes on streets in a metropolitan area. As for the propagation model, we adopt the *Two-Ray Ground Model* which has been shown [14] to provide a more accurate model of network collisions at a long distance than the free-space model. Finally, we choose IEEE 802.11 as the MAC-layer (using the corresponding implementation available in the J-SIM wireless package), since it represents the most common technology available on the market.

We set the transmission range of every node to $200m$ and run our simulations in a $4\ km^2$ area with different densities, representing respectively *sparse* (50 nodes / km^2) and *dense* (200 nodes / km^2) networks. Results obtained with various node speeds (from $5\ m/s$ to $20\ m/s$) showed no significant dependence on this parameter. As a result, we set the speed of every node to a constant value of $10m/s$.

In the following, we concentrate on a scenario with a function exhibiting a single minimum, although our approach enjoys wider applicability in multi-minima scenarios. In our simulation, each node is a potential source and publishes a new message every second. Moreover each message is associated with a propagation function defined as:

$$f(x,y) = -\frac{100000}{100000 + (x-x_0)^2 + (y-y_0)^2}$$

We set the value of v_{th} to 0.95 so that the *target zone*, A_{tg} corresponds to a circle with a radius $r = 100m$, centered around $m = (x_0, y_0)$, which in our simulation is located at $(1800, 1800)$. Notice how the target area is the same for all messages originated in the system. This generates collisions and contentions of the wireless medium close to the target area, which stresses the protocols' performance.

Dense Networks. Here, we analyze the performance of our protocols in a scenario characterized by high node density (200 nodes/km^2). This is a typical scenario representing a metropolitan area with a large number of vehicles. In this scenario, we can safely assume that the network is always connected. Hence, we focus only on the OZF, DDPD and FDPD protocols and we do not consider store & forward techniques which will be discussed next. In our analysis, we also include the *flooding* protocol in which each node re-broadcasts all the messages received for the first time. This protocol is clearly inefficient for the considered scenario; nevertheless it serves as a further baseline to evaluate the tradeoffs of our approach.

Performance results are reported in Figure 4. As expected, the flooding protocol and OZF guarantee high *delivery-IN* rates (respectively 98.4% and 99.4%)[2]. However, as shown in Figure 4(b), the number of messages forwarded by the flooding protocol is three times greater than that of OZF (135,847 messages against 41,361). Furthermore, as expected, the *delivery-OUT* of OZF is sensibly lower than that of flooding since OZF avoids backward propagation by preventing nodes from rebroadcasting a message when the the value of f at the receiver's location is greater than that at sender's.

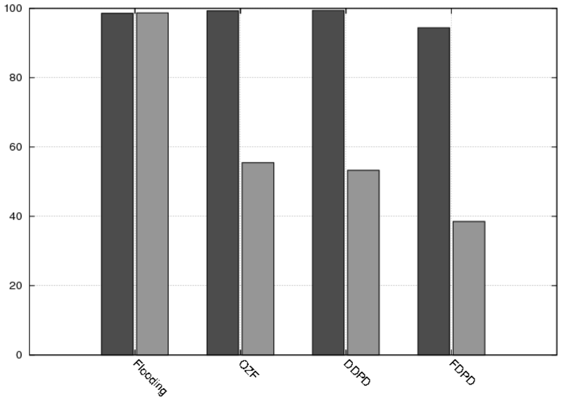

(a) Message delivery: *delivery-IN* (dark bars) and *delivery-OUT* (light bars).

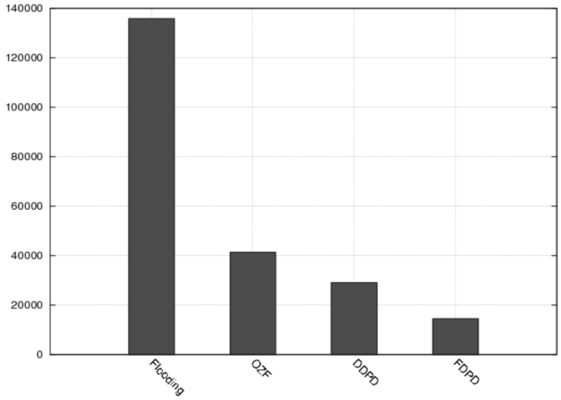

(b) Network traffic.

Figure 4: Message delivery and network traffic in a dense network (200 nodes/km^2).

Taking into account the distance, as done by DDPD, reduces the overall traffic (29,115 messages in DDPD against 41,361 in OZF) while achieving the same results in terms of delivery both inside and outside the target zone. However, both OZF and DDPD fail to provide the desired selectivity. This aspect is shown in Figure 5(a), where it is easy to see that message propagation is not optimally scoped but, instead, it spans over the vast majority of the area.

To improve selectivity and limit message propagation, we exploit the FDPD protocol described in Section 4.2 which leverages off the estimation of the best point to select the forwarding nodes. This protocol slightly decreases the delivery rate (94.4% against 99.4% in DDPD) since it forwards fewer messages at each round, but, for the same reason, it also halves the cost in terms of network traffic with respect to DDPD (14,509 messages against 29,115). This significant improvement is a direct consequence of the better selectivity of FDPD. Indeed, the *delivery-OUT* also drops from a value of 53.3% in DDPD to a value of 38.5% in FDPD. This aspect is also evident in Figure 5(b), which underlines the capability of FDPD to constrain message propagation according to the shape of the function.

Sparse Networks. The protocols discussed above are not suited for sparse networks in which connectivity is not always guaranteed. Examples of this kind of the networks can be found in rural environments with a low number of vehicles. In these challenging scenarios, as depicted in Figure 6(a), even flooding is unable to achieve a high delivery-IN because the network often becomes partitioned and there

---

[2] Packet collisions prevent these two protocols from reaching a delivery-IN of 100%.

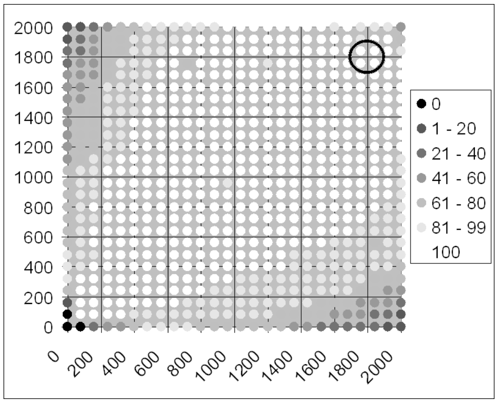

(a) DDPD.

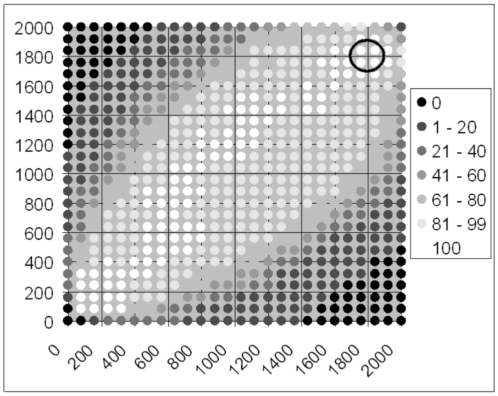

(b) FDPD.

Figure 5: Message propagation in a dense network (200 nodes/km^2).

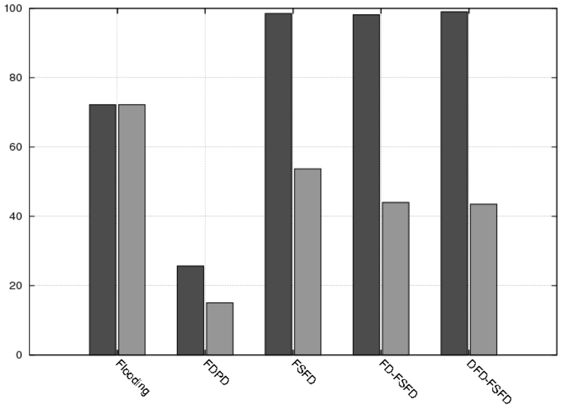

(a) Message delivery: *delivery-IN* (dark bars) and *delivery-OUT* (light bars).

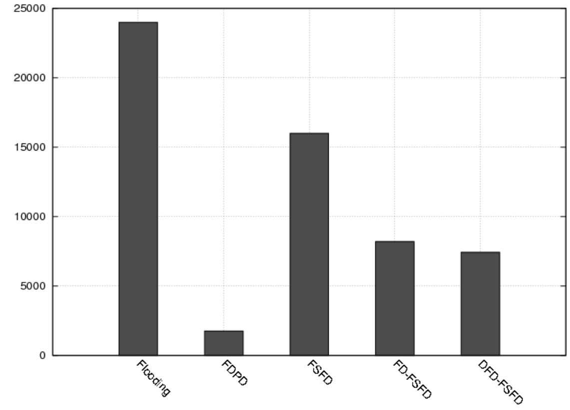

(b) Network traffic.

Figure 6: Message delivery and network traffic in a sparse network (50 nodes/km^2).

are areas in which there are no nodes that can receive and then forward a given message. For this reason, store & forward techniques become of paramount importance.

In Section 4.2, we illustrated three different protocols exploiting this kind of mechanisms. FSFD adopts the same strategy used in OZF to decide when to forward a message. In addition, after broadcasting the message, the node overhears the network channel to detect whether the message has been further propagated. This strategy yields a significant increment of *delivery-IN* which raises up to 98.5% (see Figure 6(a)) with a good *delivery-OUT* of 53.7%. Notably, despite the increased delivery, the traffic is kept low (15,997 messages against 23,990 messages in flooding). This is a consequence of the forwarding strategy and the passive feedback mechanism, which allows reliable transmission without increasing the traffic.

The FD-FSFD protocol, which integrates the FSFD protocol with the forwarding strategy exploited by FDPD, further improves selectivity and reduces the propagation area. This protocol achieves the same *delivery-IN* as FSFD but it dramatically decreases the number of messages (8,210 against 15,997) and the values of *delivery-OUT*. The selectivity improvement is also clear in Figure 7 which depicts the protocols' propagation areas.

Including the direction of movement, as specified by the DFD-FSFD protocol, improves performance only marginally: delivery shifts up to the 99% and overhead exhibits a reduction of about 10%. However, we plan to further investigate this protocol, possibly using other mobility models to provide a more accurate representation of road traffic. Such models should highlight the contribution of DFD-FSFD.

Store & forward techniques enable high delivery even in temporarily disconnected networks. However, they also result in increased latency because messages spend most of the time co-located with hosts. Figure 8 reports the average latency exhibited by our protocols. As expected, without store & forward, messages are delivered almost instantly whereas in the other cases they may take several seconds to reach the target zone. Not surprisingly, message propagation is slower in the FD-FSFD and DFD-FSFD protocols, since the probability to forward a message is lower than in the FSFD protocol.

Propagation Function. A key feature of our approach is the ability to specify arbitrary propagation functions to fit different scenarios. To verify this, we simulated the FDPD protocol in a dense network (200 nodes / km^2) using a function that forces message propagation to follow a vertical line from the lower left corner (the message source) to the upper left corner and then a horizontal line up to the upper right

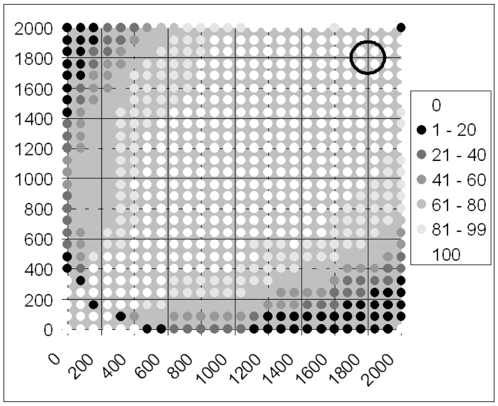

(a) FSFD.

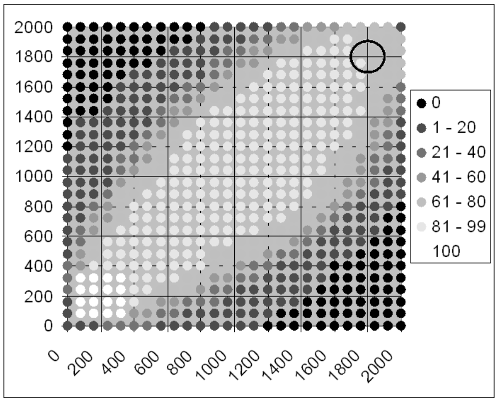

(b) FD-FSFD.

Figure 7: Message propagation in a sparse network (50 nodes/km^2).

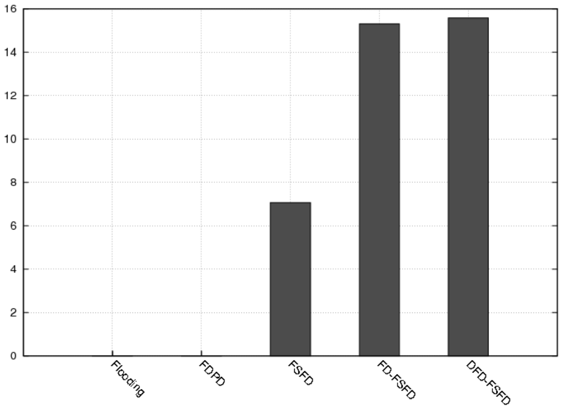

Figure 8: Average message latency in a sparse network (50 nodes/km^2).

corner, where the target area is located (see Figure 9(a)). The shape of this function is reported in Figure 9.

As it is clear from the diffusion chart in Figure 9(b), our protocol is able to propagate messages along the path imposed by the function, achieving a high *delivery-IN* together with a low *delivery-OUT*. This shows the flexibility of our approach in supporting different propagation functions that fulfill the requirements of various road environments.

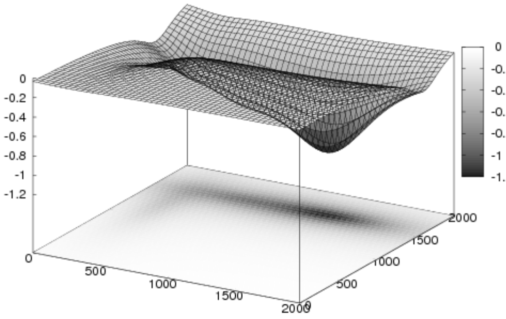

(a) Graphical representation.

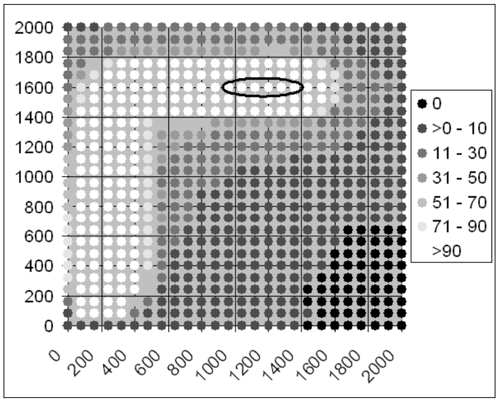

(b) Diffusion chart of the FDPD protocol in a dense network.

Figure 9: Simulation results using an L-shaped propagation function.

6. RELATED WORK

Nowadays more and more cars are being equipped with GPS receivers that enable easy retrieval of location information. This has led a significant amount of research to focus on location-based routing schemes for vehicular ad-hoc networks. Location information is commonly used both in unicast [4] and in multicast communication under the form of Geocast routing [12], the latter being a generalization of unicast location-based routing in which messages must be delivered to every node in a given geographical region.

Keeping routing efficient while avoiding obstacles such as buildings or connectivity holes is one of the most difficult issues in location-based [5] and Geocast routing [6]. Fortunately, in vehicular settings, the availability of navigation systems makes it possible to exploit map and traffic information to guide the diffusion of messages. Recent approaches examine this information to "plan" the best route to reach the destination and then use source- or trajectory-based routing [13] to diffuse messages along the desired trajectory. For example, the work in [11] computes the sequence of junctions that must be traversed by each packet to reach its destination; this information is then included in the packet in the form of geographic source routing. Another recent example of vehicular routing that exploits the availability of map information is in [16]. Its routing protocol, aimed at sparsely connected vehicular networks, uses a store and forward technique and approaches the destination by selecting the direction with the lowest estimated delay to the destination. The forwarding algorithm selects the next hop by choosing either the neighbor that is nearest to the to destination (which may lead to routing loops), or a neighbor that is approaching the target location. Different from our approach, most of these protocols address communication towards a single infostation, location, and/or connected target area, while many Vanet applications are multicast in nature (e.g., incident response coordination, congestion warnings, commercial advertisements, on-travel reservation). Moreover, while some existing work addresses the delivery of messages to multiple areas [3, 8], it does not exploit available map information as can be achieved with our message propagation function.

A propagation function neither imposes a single destination, nor selects a single path (which may fail) to follow but it instead provides a topological surface that guides probabilistic forwarding towards the target areas. To the best of our knowledge, this is the first time that such a topological function has been used in this way. While the work in [15] uses a spatio-temporal relevance function to limit the diffusion of advertisements, its approach is starkly different from ours. Specifically, the relevance function in [15] limits the propagation of messages to an area that includes the message originator, thus preventing advertisements from travelling for too long or too far away. Our propagation function, on the other hand, directs messages towards a set of target areas that do not include the sender along routes that depend on the underlying road network.

Our protocols are designed to take advantage of the propagation function while remaining simple and totally decentralized. In particular, a node forwards a message with a probably that increases with its distance from the sender. This is similar to what is done in [7], where the authors implement a reliable broadcast technique based on a Request to Broadcast (RTB)/Clear to Broadcast (CTB) scheme. Specifically, the CTB is transmitted only by the farthest node in the direction of the destination (i.e., the next junction). Potential forwarders emit a noise signal proportional to the distance from the sender and then immediately listen on the channel; the node that does not hear any noise from other nodes is the one responsible for sending the CTB, acknowledging the transmission of subsequent data and becoming the new source. While the idea of limiting the number of retransmissions is similar to that used in our FDPD protocol, the work in [7] is aimed at broadcast communication, i.e., messages are sent to the whole network. Moreover, the need to ensure reliability causes it to pay the additional overhead of RTB, CTB and acknowledgment messages.

A final aspect worth noting is that vehicular ad hoc networks are inherently extremely mobile. This makes it difficult for nodes to collect up-to-date neighborhood information as required by many existing solutions such as [16]. For this reason, our protocols are designed take decisions in a receiver-based fashion. This eliminates the need to collect neighborhood information, and makes our approach suitable for routing in highly mobile vehicular networks.

7. CONCLUSIONS

Data dissemination in vehicular networks requires flexible and lightweight routing protocols to support the needs of different applications and to scale up to the number of vehicles present in today's road networks. In this paper we presented a new way of specifying routing behavior using a propagation function that conveys information about both target areas and preferred routes. We also integrated the propagation function into several probabilistic routing protocols. Simulation results show that our protocols provide very good delivery ratios, while maintaining selectivity as well as efficiency in terms of network traffic. This validates the effectiveness of our approach and confirms its good performance in both dense and sparse vehicular networks.

Acknowledgments. The authors wish to thank Giorgio Mulas and Gianni Santo for their contribution to the implementation of the Manhattan Model. The work described in this paper was partially supported by the European Community under the IST-004536 RUNES project and by the European Science Foundation (ESF) under the MINEMA project.

8. REFERENCES

[1] J-Sim Web page. `http://www.j-sim.org/`.

[2] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. *Wireless Communications and Mobile Computing*, 2(5):483–502, 2002.

[3] C.-Y. Chang, C.-T. Chang, and S.-C. Tu. Obstacle-free geocasting protocols for single/multi-destination short message services in ad hoc networks. *Wirel. Netw.*, 9(2):143–155, 2003.

[4] S. Giordano, I. Stojmenovic, and L. Blazevie. Position based routing algorithms for ad hoc networks: a taxonomy. *Communications Magazine, IEEE*, 40(7):124–138, 7 2002.

[5] B. Karp and H. T. Kung. Gpsr: greedy perimeter stateless routing for wireless networks. In *MOBICOM*, pages 243–254, 2000.

[6] Y.-B. Ko and N. H. Vaidya. Geocasting in mobile ad hoc networks: Location-based multicast algorithms. In *WMCSA*, pages 101–110. IEEE Computer Society, 1999.

[7] G. Korkmaz, E. Ekici, F. Ozguner, and U. Ozguner. Urban multi-hop broadcast protocol for inter-vehicle communication systems. In *VANET '04: Proceedings of the 1st international workshop on Vehicular ad hoc networks*, 2004.

[8] S.-H. Lee and Y.-B. Ko. Geometry-driven scheme for geocast routing in mobile ad hoc networks. In *The 2006 IEEE 63rd Vehicular Technology Conference (VTC) (to appear)*, May 2006.

[9] L. Li, J. Halpern, and Z. Haas. Gossip-based ad hoc routing. In *Proc. of INFOCOM02*, 2002.

[10] Q. Li and D. Rus. Communication in disconnected ad hoc networks using message relay. *J. Parallel Distrib. Comput.*, 63(1):75–86, 2003.

[11] C. Lochert, H. Hartenstein, J. Tian, H. Fler, D. Hermann, and M. Mauve. A routing strategy for vehicular ad hoc networks in city environments. In *Proceedings of the IEEE Intelligent Vehicles Symposium 2003*, pages 156–161, Columbus, OH, USA, June 2003.

[12] C. Maihfer. A survey of geocast routing protocol. *IEEE Communications Surveys and Tutorials*, 6(2):32–42, 2004.

[13] D. Niculescu and B. Nath. Trajectory based forwarding and its applications. Technical Report DCS-TR-488, Department of Computer Science, Rutgers University, may 2002.

[14] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2001.

[15] B. Xu, A. Ouksel, and O. Wolfson. Opportunistic resource exchange in inter-vehicle ad-hoc networks. *MDM*, 00:4, 2004.

[16] J. Zhao and G. Cao. VADD: Vehicle-assisted data delivery in vehicular ad hoc networks. In *Proceedings of the* 25^{th} *Conference on Computer Communications (INFOCOM06)*, 2006.