

Persistent Content-based Information Dissemination in Hybrid Vehicular Networks

Ilias Leontiadis
University College London
United Kingdom
i.leontiadis@cs.ucl.ac.uk

Paolo Costa
Microsoft Research Cambridge
United Kingdom
costa@microsoft.com

Cecilia Mascolo
University of Cambridge
United Kingdom
cecilia.mascolo@cl.cam.ac.uk

Abstract—Content-based information dissemination has a potential number of applications in vehicular networking, including advertising, traffic and parking notifications and emergency announcements. In this paper we describe a protocol for content based information dissemination in hybrid (i.e., partially structureless) vehicular networks. The protocol allows content to “stick” to areas where vehicles need to receive it. The vehicle’s subscriptions indicate the driver’s interests about types of content and are used to filter and route information to affected vehicles. The publications, generated by other vehicles or by central servers, are first routed into the area, then continuously propagated for a specified time interval. The protocol takes advantage of both the infrastructure (i.e., wireless base stations), if this exists, and the decentralized vehicle-to-vehicle communication technologies. We evaluate our approach by simulation over a number of realistic vehicular traces based scenarios. Results show that our protocol achieves high message delivery while introducing low overhead, even in scenarios where no infrastructure is available.

I. INTRODUCTION

Vehicular networks are a peculiar class of mobile networks in which vehicles are equipped with radio interfaces and are therefore able to communicate with an infrastructure (if existing) or other vehicles in an opportunistic way. Content-based information dissemination enjoys wide applicability in these types of networks, ranging from traffic information and warnings, to parking availability, fuel prices, road conditions, and advertisements. While 3G networks can be used to offer these services, these come with a cost, both at network and hardware level¹, and with limitations in granularity and coverage. Our approach aims at integrating infrastructure with ad hoc approaches to decrease costs and increase coverage.

Unfortunately, while several protocols geared to content dissemination for mobile networks have been developed[1], [2], [3], [4], we have seen very few approaches specifically targeting vehicular networks[5], [6], [7], [8], [9], and, to the best of our knowledge, none which allows *persistent* dissemination in hybrid or completely infrastructureless scenarios.

In this work, we address the current research gap by presenting a protocol for *persistent content based dissemination in vehicular networks*. The protocol enables applications to:

- *publish* messages to geographical locations by first sending them to the relevant areas;

- *store* the messages in the relevant area (generally roads adjacent to where the specific event is happening) using a combinations of infostations (if any) and vehicles. Replicas of messages are stored in appropriate numbers to allow delivery to subscribers and, if stored on vehicles, they are hopped from one to another to allow them to stay in specific locations (homeZones);
- *deliver* the messages to *subscribers* (i.e., interested vehicles) inside the area, when these are met by any of the replica holders.

We take advantage of the information from the navigation system (NS), available on more and more vehicles. NSs provide valuable information on the *suggested route*. This information makes the mobility patterns of the vehicles more predictable and can be used to efficiently select the best carriers to forward messages to the affected areas and vehicles.

The suggested routes can also be used to extract interests in order to automatically filter only information relevant to the driver. For example, the NS can automatically subscribe to receive traffic warnings that affect the suggested route, to receive fuel prices from nearby fuel stations when the vehicle is running out of petrol, or to receive free parking notifications concerning the vehicle’s destination (*automatic subscriptions*). However, a user is also allowed to insert specific subscription interests, which are not automatically calculated, e.g., information on nearby restaurants or hotels (*custom subscriptions*). The subscriptions and the navigation system will then be used to filter incoming messages that concern the vehicle and its driver, to geographically route messages to/from infostations and to efficiently and persistently disseminate information in specific geographic areas. This enables our protocol to disseminate information mainly to subscribers without affecting non-interested vehicles.

We evaluate the protocol in various scenarios, using realistic traffic traces generated by various traffic simulators [10], [11], under different conditions. The results exhibit good performance in various settings in terms of overhead and message delivery, largely outperforming existing epidemic dissemination protocols [12].

II. SCENARIO

We assume a network of vehicles, equipped with navigation systems that contain information about the geographical loca-

¹The cost of a cellular chipset is currently 5-10 times higher than Bluetooth and WiFi chips.

tion of local infostations (i.e., access points to the backbone) and the planned route to reach the desired destination. Each vehicle is provided with an omnidirectional antenna and is able to wirelessly communicate with neighbouring vehicles. In addition, we assume that power consumption is not critical and that storage space is virtually unlimited: this is reasonable for vehicular networks.

Information is generated by *publishers* in the network: these can be either central servers on the backbone or vehicles themselves. The publisher defines the area where the information should be disseminated (*Persistence Area*). Publications can be the result of information collected beforehand in the network or fresh information only distributed locally by vehicles directly. We assume information is collected by a centralized system: it is not the scope of this paper to discuss how it is collected, we can even assume that the very same vehicles are mobile sensors that collect information about traffic conditions, accidents, etc (e.g., like in CarTel [13]).

A vehicle, may act as a subscriber, by expressing its interests in a certain set of messages. In vehicular applications, relevant information ranges from traffic news (e.g., road works or congestion) to gas stations and hotel advertisements. Figure 1 describes one of our scenarios: a car needs to travel on a certain road, indicated with the red arrow following the main yellow road. At a certain point, road works are scheduled or an accident happens on that road (indicated in the right side of the picture): we will call this *Point of Interest* (POI) from now on. The cars heading in the directions of the POI need to be informed so that they can potentially follow another route.

As illustrated in Figure 1, our approach aims at letting information *stick* to an area (generally the area leading to the POI) for a certain time period. This is done by creating *homeZones*, which are locations where message replicas are stored for distribution to subscribers. The homeZones will be placed on the road segments where subscribers drive to reach the POI. As we will see later, the homeZones represent either the location of road-side infostations, if present, or the location where the replicas should be kept by vehicles traveling nearby.

In order to keep disseminating information, we use content based dissemination which allows us not to inform all vehicles in the area but just those interested. Content-based routing (CBR) differs from classical routing paradigms as messages (i.e., the published information) are routed based on their *content* rather than their destination address. This form of implicit, multi-point communication fosters a high degree of decoupling, since the communicating parties are not necessarily aware of each other, and can therefore change dynamically without affecting the rest of the system.

III. PROTOCOL DESCRIPTION

Given the heterogeneity of the scenarios we target, a widespread presence of infostations cannot be guaranteed at any time and any place. Therefore, we developed a novel approach, striking a balance between the efficiency of infrastructure and the flexibility of opportunistic communication.

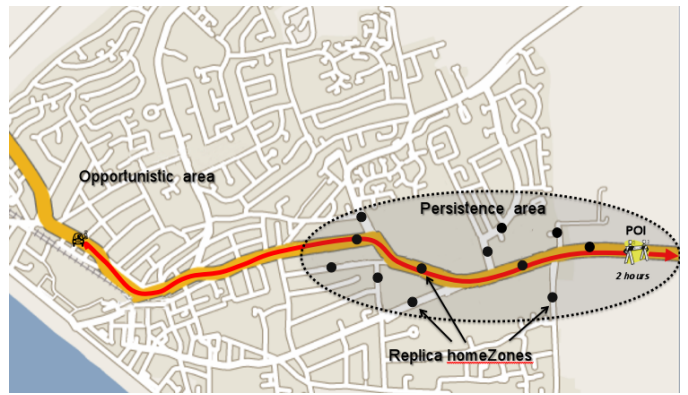


Fig. 1. Black dots represent homeZones where messages should be stored to inform approaching vehicles about the road works (POI).

Indeed, as we detailed in the following, our protocol seamlessly exploits both infostations (where available) and vehicle-to-vehicle communication to *deliver* and *store* messages in the intended locations.

Hereafter, for sake of clarity, we illustrate our approach in different steps by first describing the basic version of the protocol, assuming pervasive infrastructure availability, and then we show how we can relax this assumption, incorporating opportunistic and ad hoc communication to i) extend the infostation dissemination range and ii) maintain the message persistence even in areas without infostations. Nevertheless, these three variants are not independent but co-exist in our protocol to provide a single solution addressing content-based dissemination in heterogeneous environments.

A. Infrastructure-based Persistence

As detailed in Section II, the aim of our protocol is to ensure that all drivers are promptly informed about relevant events, affecting their route (e.g., traffic jams or gas stations). To this end, information about these events must be stored at specific locations, called *homeZones* (indicated by black dots in Figure 1) such that all approaching vehicles can be notified.

The identification of the exact position and the number of these homeZones can be done automatically or in an application specific way and will vary according to the type of information and the road topology. For example, we can use a simple algorithm to make sure that there are replicas in every path leading to the POI from a certain distance. Or, a highway agency can strategically define the areas where the information should be persistent: for instance, in case of the traffic jam on a highway in Figure 1, the homeZones sit on the main roads to access the highway so that vehicles can avoid entering the highway and choose alternative paths.

For each homeZone a *replica* of the original message is created and then routed and stored at the corresponding geographic location. Under the assumption of widespread infrastructure, messages can be sent to the nearby infostations to be then disseminated. A simple yet inefficient solution would be to have each infostation to periodically re-broadcast

the message to all nearby vehicles. This, however, would incur a significant network overhead because i) messages are transmitted even if no subscribers are around and ii) subscribers are very likely to receive the same message multiple times by encountering several infostations along their paths.

To circumvent this issue and to remove unnecessary transmissions, we devised a two-phase scheme where each vehicle periodically advertises its planned route and its additional interests, e.g., fuel or parking slot, if any. Through this information, the infostations can derive the actual subscriptions (both automatic and custom) and compare them against the stored messages.

We assume that subscriptions and messages are matched through a **match** function. Messages are defined as a list of attributes values (e.g., `<gas company="abc", fuel type="unleaded", price per liter="1.25">`) whereas subscriptions are expressed as queries over these attributes (e.g., `<*, diesel, less than 2>`) but more sophisticated patterns including regular expressions can also be used.

If a match occurs, the message is then transmitted and received by the subscribers. To avoid duplicate receipts, the subscriber also piggybacks the IDs of the last λ messages received. In this way, before forwarding the message, the infostation could check whether that message has already been delivered. Only subscribers that did not previously receive the message can trigger a broadcast, which, however, can be heard by more than one subscribers in the area.

B. Opportunistic Dissemination

Even in presence of infostations, opportunistic vehicle-to-vehicle communication can greatly enhance the performance of the above protocol by enabling the dissemination of messages in a broader area at a very small additional cost. To this end, we let the vehicles (subscribers and non) which have heard a message store it and retransmit it when detecting a subscriber in the vicinity. This allows for opportunistic exploitation of vehicles which in any case have heard the information and can act as additional carriers for the information in the persistence area and beyond it when vehicles exit it. Interestingly, if no subscriber is encountered, no additional traffic is generated, thus effectively implementing an *interest-driven* routing scheme in which messages propagate only in areas populated by subscribers, possibly extending the persistence area defined by the application.

C. Ad-hoc Persistence

In some scenarios the existence of infostations is not sufficient to allow the dissemination to all interested vehicles. This could be for various reasons: i) the infrastructure is partially collapsed due to accidents or attacks ii) the infrastructure is not covering the whole interested area as this is vast iii) the information has a very fine granularity with respect to the infostation coverage (e.g., parking slots positions may be of interest only to be persistently maintained only in a couple of streets away from the parking area).

With respect to the approach presented in Section III-A, if infostations are not available, we then need to find ways to i) *route* message from the publisher to the homeZones and to ii) *keep a replica* persistently close to its homeZone.

While the just described opportunistic dissemination helps to partially solve these issues in a hybrid scenario, it does not guarantee that a replica is persistently maintained in the homeZone as most of the vehicles which have heard the message might leave. Alternative approaches based on 3G technology like UMTS are not viable as the available bandwidth per user in a crowded area would be severely limited and the price the users should pay to network companies would discourage them from adopting the service.

Therefore, a more proactive mechanism is needed to enforce persistence in a semi or totally decentralized scenario, relying on *scalable* and *inexpensive* ad hoc communication among vehicles and exploiting the route information provided by navigation systems.

While power consumption is not critical in vehicular networks, bandwidth is a critical factor which may drastically reduce the capacity of wireless communication [14] and limiting its usage it is of utmost importance to enable ad hoc communication. It is therefore unfeasible to store messages on all vehicles and let them flood the network because this would rapidly saturate all the available bandwidth. Hence, it is key to bound the number of transmissions by properly identifying which are the nodes interested in a given messages and which are the best carriers to keep the replica close to its homeZone. Ideally a carrier would be a vehicle which is always near the replica's homeZone. This is clearly unrealistic as in general vehicles continuously move from one location to another and, hence, new carriers must be selected. In particular, not only the current location of a node is important but also its future one as carriers moving *towards* a homeZone are much better than those departing *from* it.

Notably, however, while in traditional mobile networks predicting user movements is hard (if not impossible) in vehicular networks information about future movements can be derived by looking at the planned route provided by navigation systems. This is particularly beneficial as a host can detect whether one of its neighbors is a better carrier for some of the buffered messages. Note that information about future route is broadcasted even in the infrastructure-based version of our protocol to derive automatic subscription so no additional traffic is required.

We introduce the notion of *utility* for the selection of message carriers. The utility $U_v(m)$ of a vehicle v with respect to a message m represents how good of a carrier v is to deliver m to its homeZone. This depends on how close that node will be w.r.t. the homeZone, i.e., is the nearest point (NP), and how fast the node will get to there according to its travel speed. Then, we apply the Dijkstra algorithm to estimate the minimum time occurring to reach the homeZone under the assumption that a new carrier is found.

More formally, we have:

$$U_v(m) = T_{NP} + \hat{T}_{homeZone} \quad (1)$$

being T_{NP} the time needed to reach NP given the current route and speed of the vehicle and $\hat{T}_{homeZone}$ the estimated time to go from NP to the homeZone based on the Dijkstra algorithm².

For instance, in the example depicted in Figure 2, a message is published at the infostation I and needs to be route to the homeZone denoted by the red dot. Both the vehicle V_A and vehicle V_B are potentially eligible as carrier but the latter is preferable because its route will get closer to the final destination, i.e., the message’s homeZone. Nevertheless, while approaching NP_B the node encounters another vehicle, V_C , whose route happens to cross the homeZone and hence it takes over the message and delivers it to the destination. Note that this process never ceases because as soon as V_C will pass the homeZone, it will need to find another carrier going in the opposite direction back to the homeZone.

Thus far, we focused our attention only on carriers but the goal of the protocol is to deliver the message to the subscribers. To this end, when a node receives the planned route from its neighbors, beside checking whether there is any potentially better carrier, it also verifies whether its messages are of interest for any of its neighbors, adopting the same approach as in the infrastructure-based version (e.g. replica routing and information dissemination occur at the same time). Therefore, we distinguish between two kind of messages: *replicas* and *notifications*. The former are the messages that need to be located as close as possible to the corresponding homeZone and can never be deleted. Notifications, instead, are those messages that have been delivered to subscribers or overheard.

D. Discussion

Our approach largely relies on routes available in the navigation systems. This implicitly assumes that users are cooperative and willing to insert their destination. One might argue, however, that this assumption is only partly verified in practice as users tend to avoid using navigation systems for known routes (e.g., when going to the work places). Nevertheless, we expect the drivers will have an incentive to insert their destination in the navigation system as this action will automatically subscribe them to any type of notification about it, thus being able of receiving critical information about traffic congestion, accidents, etc which are of utmost importance even (if not more) for daily routes.

Similarly, at a first glance, privacy concerns seem to hamper the adoption of our protocol since drivers may be reluctant to publicly advertise their route. While we do not feel this being an issue for most, it can be easily solved by employing specific techniques to anonymise this information. Although a throughout discussion of these aspects is out of the scope of the paper, a simple yet efficient mechanism would be to

²Note that this calculation can be easily executed by the on-board computer by leveraging off the map and the planned route of the neighbor v .



Fig. 2. Routing a message replica to its homeZone

use encrypted transmissions and to frequently change the ID associated to a given vehicle³.

IV. EVALUATION

To evaluate our approach, we report on the protocol performance over several simulated scenarios, generated from realistic vehicular traces. We analyzed our protocol under a synthetic load of both automatic and custom subscriptions. In particular, for automatic subscriptions, all vehicles with planned route intersecting the POI are considered subscribers. This is the typical situation with traffic warnings, which are of interest to any vehicle enroute towards the affected destination. Conversely, custom subscriptions (e.g., hotel or restaurants) are not relevant for everybody but will involve only a fraction of vehicles travelling towards the POI.

To put our work in the context of related efforts and to capture the tradeoffs involved, we compared our solution with an epidemic approach, reminiscent of [12], in which all nodes store each message received and re-broadcast it to all neighbours, which have not heard that message yet.

A. Simulation Settings

As simulation platform, we used *OMNet++* [15], an open-source discrete event simulator and the *mobility framework plug-in* [16]. We use 802.11b wireless radio interface (max range is 250m). In the default configuration we have 450 vehicles, advertise interval equal to 10 s and 10 replicas. Each simulation lasts for 2 hours of simulated time and results are averaged over multiple runs.

To accurately assess the performance of our protocol in the context of vehicular networking, we exploited traffic traces generated by a multi-agent microscopic traffic simulator at ETH, Zurich [10]. These traces contain mobility patterns of 260,000 vehicles over real road maps in the canton of Zurich within a period of 24 hours. However, we extracted smaller areas (50x50 km) to decrease the duration of the simulation (by implicitly reducing the number of vehicles involved). Additionally, we used the GMSF generator [11] to produce

³Given the high mobility, ID collisions are unlikely to happen and, hence, IDs can change often without harnessing the correct behavior of the protocol.

GIS traffic-light traces for the rural, urban and city scenarios, which have finer granularity (3x3km).

The characteristics of the scenarios considered are:

- *City scenario*: High vehicle and street density scenario where up to 880 vehicles are concurrently present (default 700). We place the POI on an intersection of two secondary roads: subscribers can be up to 40 of the 700 vehicles (automatic subscriptions). Average speed near POI is 20km/h and maximum is 60km/h. An example of this is provided in Figure 7(a) where POI and replicas are also shown.
- *Urban scenario*: Medium street and vehicle density. 420 Vehicles are present at the same time and up to 30 of them can be subscribers. Maximum speed is 60km/h but average speed is 25km/h (Figure 7(e)).
- *Rural scenario*: This is a low density scenario where only 100 vehicles concurrently present. In the scenario of Figure 7(i). Subscribers can be up to 40 of the 100 vehicles (automatic subscriptions). Average speed is 28km/h (max is 60km/h).
- *Highway scenario*: This is a much larger 50x50km area as illustrated in Figure 7(m). The simulation includes a maximum of 830 concurrent vehicles and up to 350 can be subscribers (most of the vehicles will drive near the POI). The average speed is higher than the previous scenarios (93km/h) and the highest 120km/h.

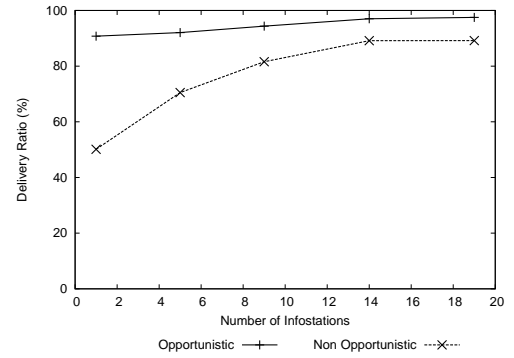
B. Simulation Results

Hereafter, we will first present results achieved in the city-base scenario, with and without infostations, as this represents the more challenging case for our protocol, given the complex road topology. Then, we will show the performance obtained in the urban, rural, and highway scenarios to demonstrate the suitability of our approach to different environments.

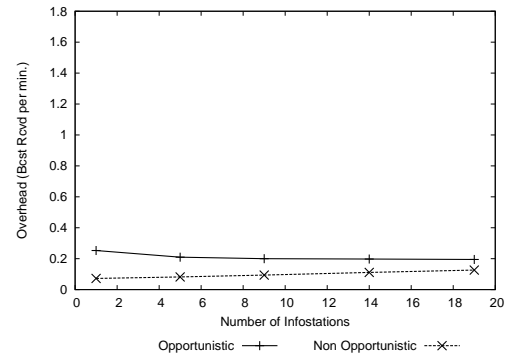
In all our experiments, we measured the *delivery ratio*, expressed as the fraction of subscriber that successfully received the messages; and the *network overhead*, defined as the number of transmissions received per minute by each vehicle.

Infostations. As a first experiment, we focus on a fully infrastructure-based scenario in which the persistence area is instrumented with several infostations. Our goal is twofold: on one hand we want to demonstrate the correctness of our protocol and on the other hand we want to assess the impact of the additional opportunistic dissemination in such scenario. To this end, in Figure 3(a) we measured the delivery of our protocol under two different configurations, i.e., with and without opportunistic dissemination.

Remarkably, through the opportunistic dissemination introduced in Section III-B, delivery is above the 90% even with just one infostation. On the other hand, if opportunistic dissemination is not used, at least 14 infostations are needed to achieve similar performance. This is a prominent result as it proves that even in a fully infrastructure environment, opportunistic dissemination represents an asset to our approach. Indeed, although the network overhead does not change with the number of infostations (see Figure 3(b)), still resorting



(a) Delivery Ratio (high density).



(b) Overhead (high density).

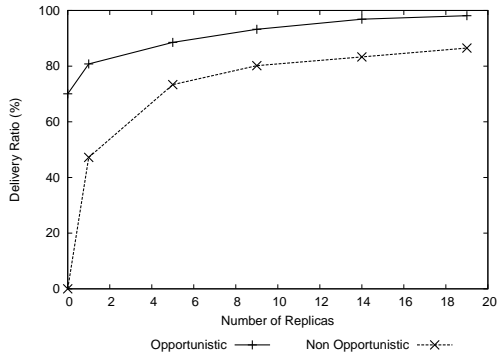
Fig. 3. Number of Infostations.

to opportunistic dissemination enables the reduction of the number of infostations, thus simplifying their deployment.

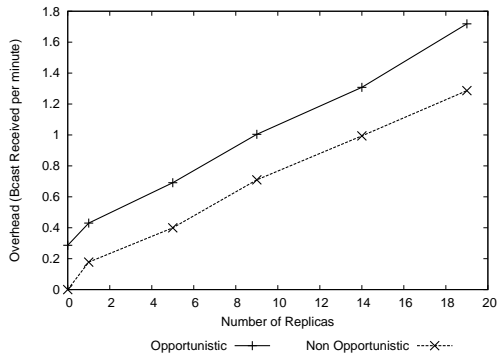
Ad-hoc. Nonetheless, despite the above results, assuming a widespread availability of infostations is unrealistic in many scenarios. Hence, to ensure efficient content-based dissemination in hybrid scenarios, as those targetted in this paper, it is fundamental to support infrastructure-less communication. In our work, this is achieved by means of the *ad hoc* persistence solution, described in Section III-C. To avoid any bias and to isolate the contribution, in the rest of this section we will assume that no infostation is present and that all communication relies on vehicle-to-vehicle technology. Clearly, in case of semi-infrastructure environments, we can have an interplay of the two approaches.

Number of Replicas. The first parameter we explore is the number of replicas created to guarantee the persistence of the message within the specified area. Also, as we did in the infostation scenario, to assess the impact of the opportunistic dissemination, we run two different versions of our protocol: the former relying only on replicas to disseminate messages and the latter exploiting also the opportunistic routing. Since results strongly depend on the density of vehicles, we tested it both in a low and high density scenario (200 and 700 vehicles).

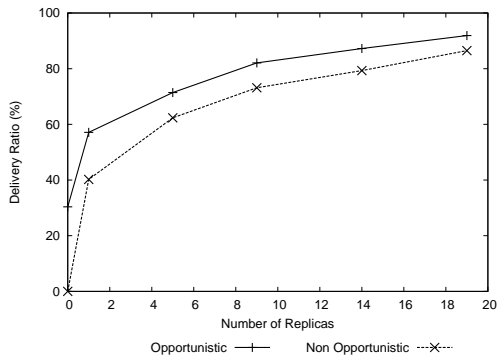
Results in Figure 4(a) confirm our claims. When the density is high, even a small number of replicas is sufficient to achieve a high delivery. Interestingly, however, this remarkable result



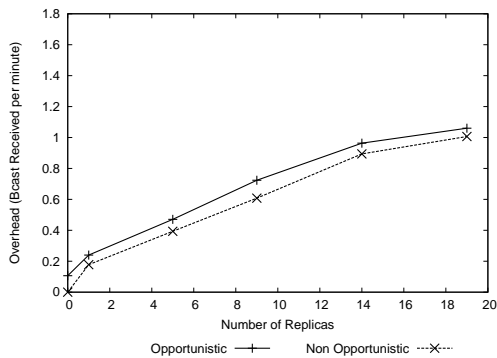
(a) Delivery Ratio (high density).



(b) Overhead (high density).



(c) Delivery Ratio (low density).



(d) Overhead (low density).

Fig. 4. Number of Replicas.

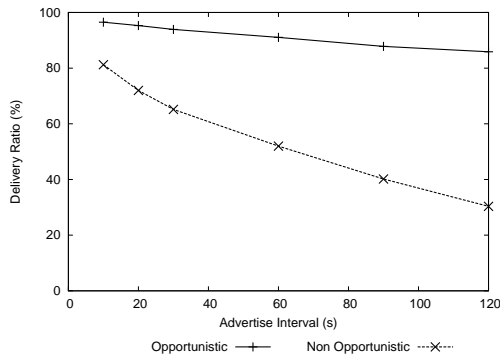
is due to the combination of two different strategies: the ad hoc persistence and the opportunistic dissemination. Indeed, when the opportunistic dissemination is not used, the delivery drops to 50%, unless many more replicas are introduced. This however, as shown in Figure and 4(b), generates a significant overhead. Indeed, to achieve the same delivery of 80%, 9 replicas are need without opportunistic dissemination (instead of just 1) with almost doubled overhead (0.7 against 0.4 broadcasts per minute). Notably, the opportunistic dissemination only slightly affects the overhead because the most of it is due to keep replicas in the persistence area. Furthermore, if opportunistic dissemination is not used, even a high number of replicas does not bring significant improvements to the delivery.

In case of low density, as expected, the overall improvement provided by the opportunistic dissemination decreases as there are fewer vehicles around. Hence, the main transmissions will occur from replica carriers and this explains why the delivery is mainly impacted by the number of replicas. Nevertheless, the opportunistic dissemination is still useful because it yields an improvement of about 10% in terms of delivery regardless how many replicas are used.

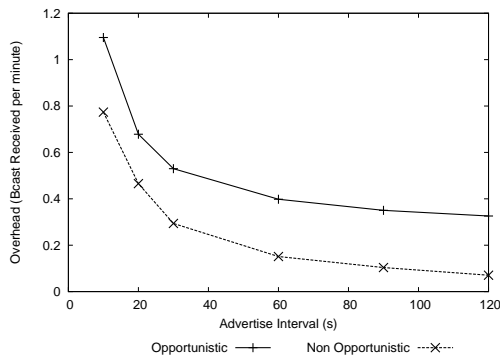
Looking at these results, one might argue that the main contribution to the message delivery comes from the opportunistic dissemination while the ad hoc persistence plays only a marginal role. This, however, is strongly contradicted by performance achieved with zero replicas, both in the high density and, especially, in the low density scenario. Indeed, in the former, opportunistic dissemination alone delivers the message only to the 70% of subscribers while in the low density scenarios only the 30% of subscribers are notified. This is consistent with the conclusions drawn above: opportunistic dissemination provides a valuable contribution only in dense scenarios while in sparse scenarios it becomes less useful. Nevertheless, even in dense networks, to get reasonable results, it must be coupled with persistence strategy since, otherwise, if the message disappears from the area, by no means later subscribers can be notified.

The results in the high density scenario (Figure 4(a)) closely resemble the ones with infostations in Figure 3(a). Not surprisingly, however, overall performance is slightly worse: This behavior stems from the fact that now replicas are hosted on vehicles, as opposed to infostations. Hence, even non-subscribers play a key role to ensure proper persistence, by continuously passing replicas from one vehicle to another. Delivery of subscribers is also affected because in some cases, replicas may abandon the homeZones (e.g., because no alternative carriers were found). Consequently, incoming subscribers may miss the notification, thus demanding for more replicas to be in place.

Advertise Interval. Advertise interval is a complementary parameter w.r.t. the number of replicas. If we keep the number of replicas fixed, we can reduce the advertise interval to improve the message delivery. In this way, the probability for a subscriber to miss a replica is lower because subscribers advertise their interests more frequently.



(a) Delivery Ratio (high density).



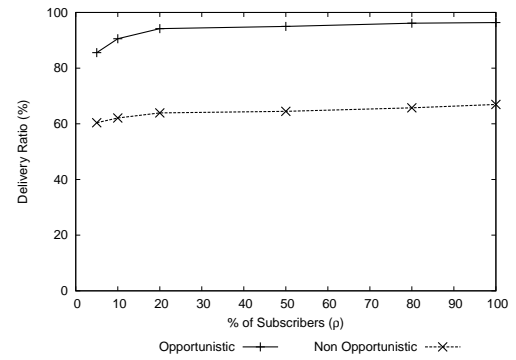
(b) Overhead (high density).

Fig. 5. Advertise Interval.

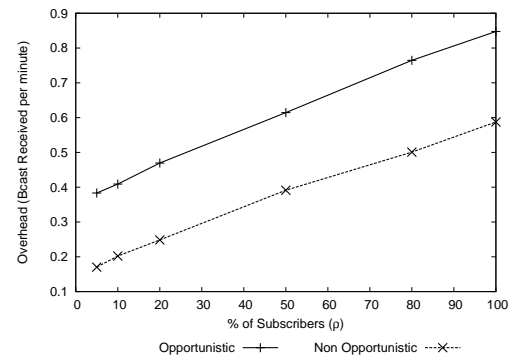
This property is charted in Figure 5 in which we studied the protocol behavior over different advertise intervals. As described above, decreasing the advertise interval is beneficial to the delivery which increases to almost 100% (here we used 10 replicas). Interestingly, the improvement in terms of delivery is more evident when opportunistic dissemination is not used: without opportunistic dissemination, missing a replica is far more critical because the chances to encounter another one are few. Conversely, opportunistic dissemination alleviates this issue since messages can be obtained also from other vehicles and not exclusively from replica carries.

Note, however, that reducing the advertise interval comes at a cost. Beside incrementing the advertisements per minute, it increments the overall number of broadcasts received. Indeed, given that information about nearby vehicles is more accurate, replicas will hop more frequently from one vehicle to another because better carriers are found. This explains why the number of broadcast exhibits a steep trend as soon as the advertise interval gets small.

Custom Subscriptions. Thus far, we concentrated our attention only on automatic subscriptions. Nevertheless, a prominent feature of our approach is the ability to incorporate also driver's interests, which are not necessarily shared by all other drivers. To model this scenario, we assume that only a fraction ρ of vehicles going towards the POI are actually interested in the message and we analyze our protocol under different



(a) Delivery Ratio (high density).



(b) Overhead (high density).

Fig. 6. Custom Subscriptions.

values of ρ (see Figure 6).

Remarkably, as reported in Figure 6(a) our protocol shows high event delivery even for small values of ρ . This means that regardless of the fraction of subscribers, our protocol ensures that the vast majority (e.g., 90% for $\rho = 10\%$) of them receives the message. Furthermore, we also observe that when there are more subscribers, the message overhead increases. This verifies that low interest messages are spread less than more popular ones (i.e., the spread/overhead depends on the interest about an event).

These charts demonstrate the high flexibility of our protocol, which is able to tune to network conditions and to selectively contact almost only intended subscribers.

Distribution. In all previous charts, we focused our analysis on the city scenario, since this represented the most challenging test for us. Nevertheless, to carefully evaluate our protocol, we experimented also with other traces, available at [11], representative, of an urban, a highway, and a rural scenario and compare it with results obtained in the city scenario.

We first plot the distribution of informed vehicles to get a visual intuition of the performance of our protocol in the three scenarios, as depicted in Figure 7. Looking at the Figure 7(a), 7(e), 7(i), and 7(m), the different topologies of the four scenarios emerge. In the city scenario, much more roads and potential routes are present while in the latter three, the topology is simpler. Figure 7(b), 7(f), 7(j), and 7(n) depict the distribution of subscribers across the whole simulation area.

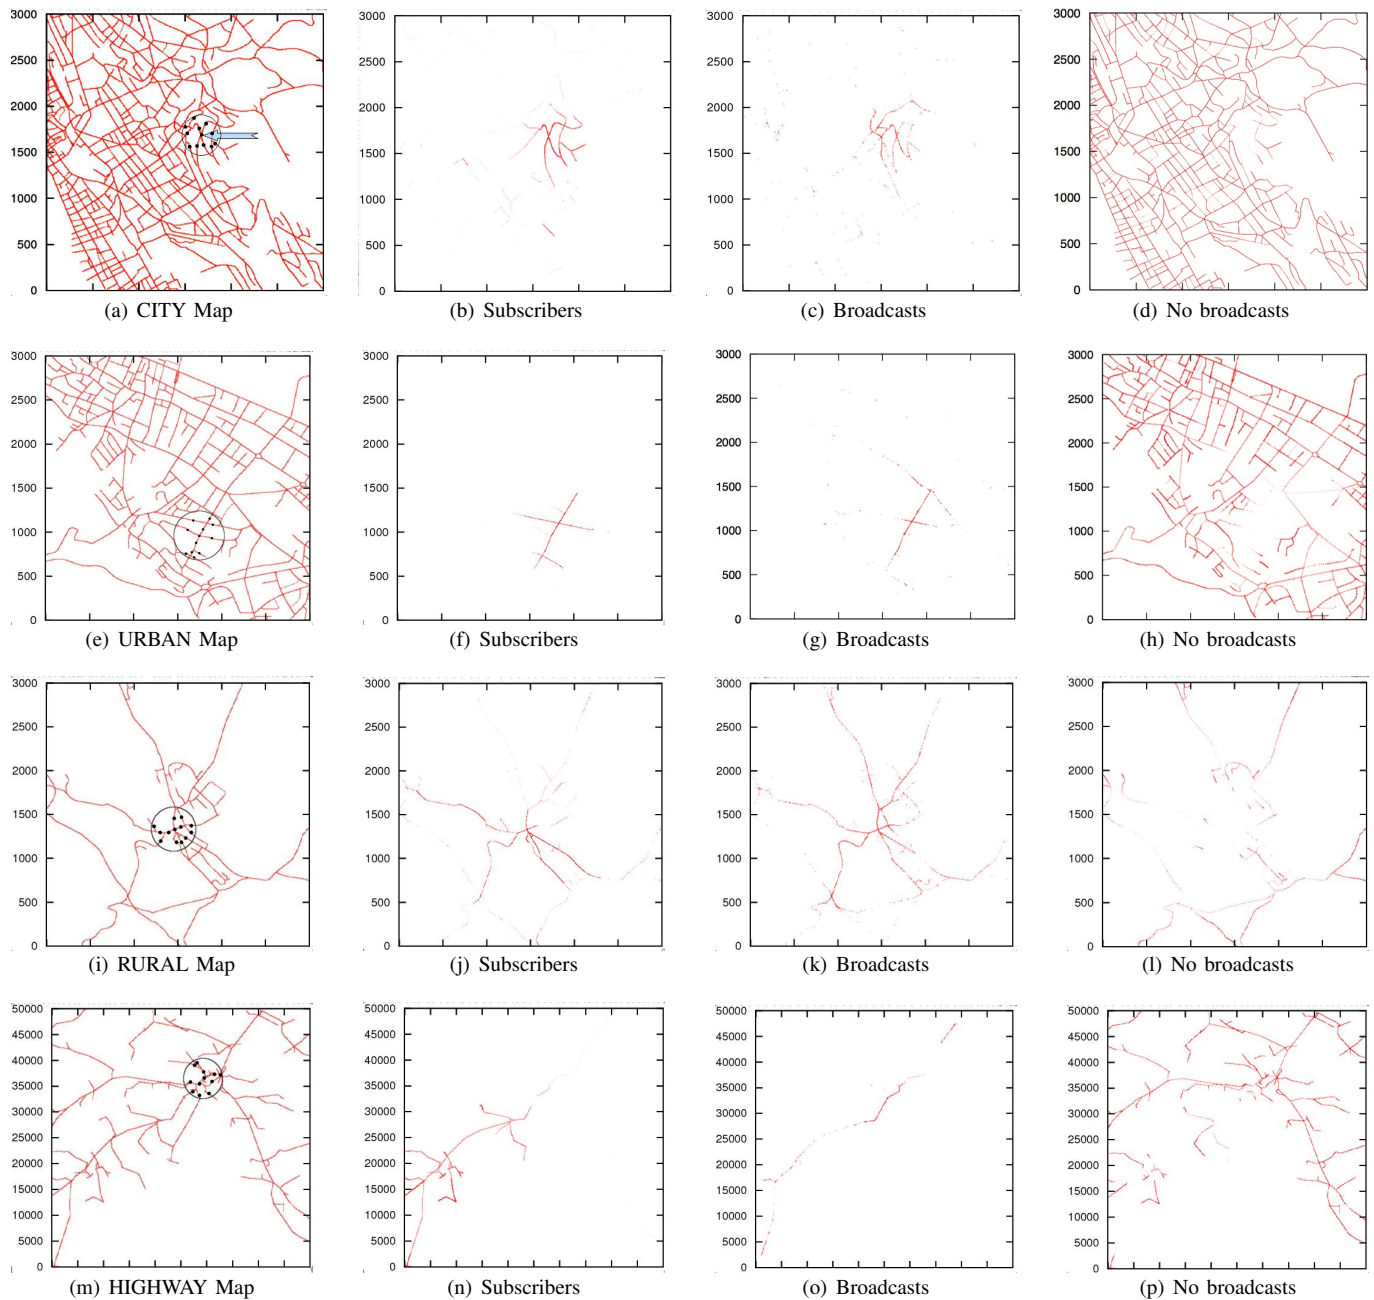


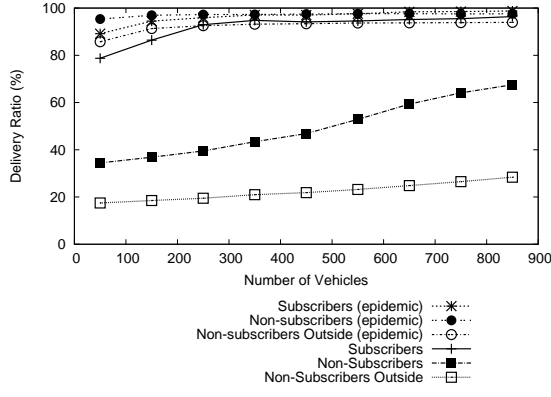
Fig. 7. City (a-d), urban (e-h), rural (i-l), and highway (m-p) scenarios. First column illustrates the Map, POI, replicas (black dots), and persistence zone (circle). Second contains road segments with high percentage of subscribers. Third depicts the broadcast distribution while the fourth demonstrate road segments with no broadcasts.

Note that these include all nodes travelling towards the POI depicted in the left most charts. Because of the more complex topology, in the city scenario, only nodes close to the POI are actually subscribers while in the other scenarios, since there are fewer roads, all nodes travelling on the main road are subscribers, i.e., all nodes are going towards the POI.

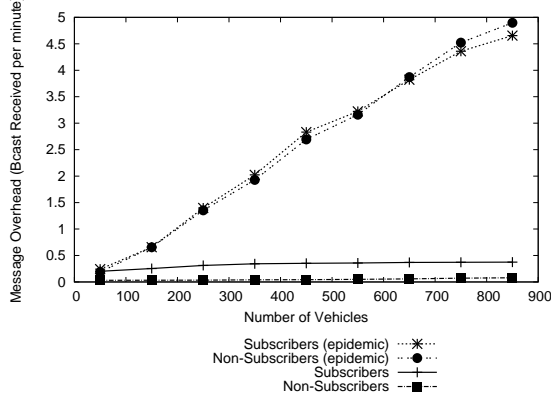
Regardless the underlying topology, the main contribution from the delivery, as already outlined, comes from the replicas in the persistence area. Indeed, the distribution of broadcasts (see Figure 7(c), 7(g), 7(k), and 7(o)) is higher in the

persistence area than in the rest of the chart, as plot in Figure 7(d), 7(h), 7(l), and 7(p).

Note that message propagation extends also beyond the persistence area but almost only subscribers are reached by the message. This behavior is due to the opportunistic dissemination which keeps on informing new subscribers, exploiting vehicles which overheard the message in the persistence area. In this way, subscribers are informed, at virtually no cost, much earlier than the time they would enter the persistence area, thus enabling them to take the proper actions, e.g., in case

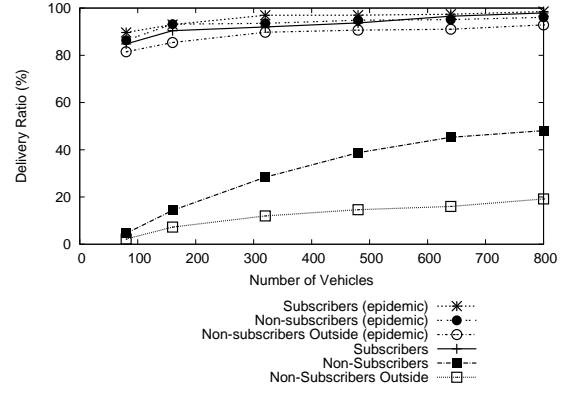


(a) Delivery inside and outside the persistence area.

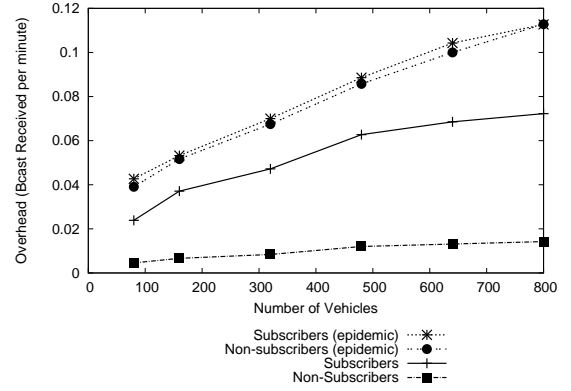


(b) Overhead

Fig. 8. Density of vehicles (City)



(a) Delivery inside and outside the persistence area.



(b) Overhead.

Fig. 9. Density of vehicles (Highway)

of a traffic congestions or emergency, in advance. This is even more evident in the highway and rural scenario because, given the scarcity of roads, subscribers leaving the persistence area are much likely to travel on the same road, but in the opposite direction, of subscriber going towards that area, thus increasing the probability of opportunistically exchange messages.

Finally, if vehicle density is low, e.g., in the rural scenario, replicas can leave the persistence area because the current carrier might not find any suitable vehicle to forward the replica and, hence, the replica is kept until a better carrier is encountered. This explains why in Figure 7(k) we have some broadcasts also in areas where there are no subscribers.

Density of Vehicles. To get further insights on the efficiency of our protocol, we compared it against an epidemic version, inspired to [12]. In this protocol, all nodes gossip to all neighbours which have not previously received the message. In this way, the epidemic infection is kept alive and eventually all vehicles gets informed. This protocol can be seen as an extension of our opportunistic dissemination in which all vehicles, not just subscribers and vehicles which overheard it, receive the message. We already shown in Figure 4 and 5 that opportunistic dissemination is not sufficient unless coupled with persistence (either infrastructure-based or ad hoc). Here we make a further step in this direction and show that

epidemics provide good performance in terms of delivery but the overhead is order of magnitude higher than ours. This is observable in Figure 8(a) and 8(b): although delivery is quite high, the overhead increases enormously. Furthermore, while the overhead of our protocol increases sub-linearly density of vehicles, the epidemic overhead increases linearly. The difference is due to the selectivity of our protocol which delivers message only to proper subscribers and hence it is less impacted by the density of vehicles. On the other hand, the epidemic protocol infects all vehicles, not just subscribers, as illustrated by the much higher delivery ratio of non-subscribers. This becomes even more critical if we extend our analysis of non-subscribers outside the persistence area. Indeed, while, outside the persistence area, our protocol affects around the 20% of vehicles non-subscribers, the epidemic protocol has to contact *all* vehicles, which is unacceptable in real situations.

The same trends are observed in the highway scenario in Figure 9(a), although here most nodes are subscribers and hence the fraction of non-subscribers informed is much lower with our approach. The overhead in Figure 9(b) follows a behaviour akin to the one observed in the city scenario, although the absolute values are lower. On the highway the set of neighbors changes less frequently and, hence, broadcasts

are less triggered. Similar tradeoffs also emerged in the urban and rural scenarios (not shown for space reasons).

These results fully confirm that our protocol deals effectively with the characteristics of hybrid vehicular networks, ensuring high event delivery with reasonable overhead in a heterogeneous set of realistic scenarios.

V. RELATED WORK

To the best of our knowledge, Abiding Geocast [3] is the most related work in terms of delivering time-stable message in a geographical area. In order to disseminate the message in the area it employs periodic flooding or epidemic dissemination. With respect to this work, our protocol considers subscriptions, topics and navigation systems in order to filter and route the message to the interested vehicles (it is a content-based dissemination protocol and not a Geocast protocol). In addition, through the use of ad hoc persistence, we drastically reduce the overhead as opposed to epidemic approaches.

LPS [1], L-ToPSS [2], and STEAM [4] are location-based publish-subscribe systems in which location is defined as a range from the publisher or subscribers. Conversely, in our approach we have detached the persistence area from the subscriber or publisher location, thus improving the flexibility. Also, none of these approaches enable persistent dissemination. They provide delivery only to nodes which are inside a destination region exactly at publishing time while instead we allow messages to persistently remain in a given area.

In the field of vehicular networks, some works exploited contextual information to steer message propagation. In [17] a system for dissemination of information in an area is presented: the framework, however, does not take advantage of topological constraints of the road map to drive the dissemination and only concentrates on notifying a percentage of nodes in a circular area. Works targeting multicast communication in vehicular networks recently appeared in literature [9], [18], [19], [20], [5]. They used different versions of scoped epidemic protocols to constrain the propagation of a message within the given area specified by the publisher. Other works [6], [8], [7], instead, define a notion of relevance to enable the routing layer to self-identify the areas in which the messages should be delivered. In contrast to these approaches, our work offers a much richer semantics in which publishers and subscribers are completely decoupled as the former define the persistent area while the latter express their interests and these are matched against their route to filter out unnecessary information.

Nonetheless, beside the aforementioned differences, the most prominent novelty of our approach is that none of the above approaches targeting vehicular networks support persistence. This is a major flaw because the vast majority of information disseminated (e.g. warnings, gas station advertisements) are usually meant to last several hours. Hence, the only way to ensure this in the cited works is either to periodically re-broadcast the message (thus largely degrading the overall performance of the system) or exploiting an epidemic-like approach in the targeted areas, which, however, as we discussed

in Section IV, would introduce unacceptable overhead.

VI. CONCLUSIONS

We have presented a protocol for persistent content dissemination in hybrid vehicular networks. Messages specify a point of interest and a persistence area (i.e., a set of roads) in which the information needs to remain in order to notify subscribers. Subscribers are vehicles interested in the information. Our mechanism exploits the navigation system information to generate and match subscriptions, to route and store the message in the relevant area. We have evaluated our approach by using realistic traffic traces. The results show good performance in various settings in terms of overhead and message delivery, also with respect to epidemic dissemination.

Acknowledgments: We would like to acknowledge the support of the EPSRC through the Project CREAM (EP/C544765/1).

REFERENCES

- [1] P. T. Eugster, B. Garbinato, and A. Holzer, "Location-based publish/subscribe," in *In Proc. of 4th Symp. on Network Computing and Applications*, 2005.
- [2] I. Burcea and H.-A. Jacobsen, "L-ToPSS - Push-Oriented Location-Based Services," in *TES*, 2003.
- [3] C. Maihofer, T. Leinmuller, and E. Schoch, "Abiding geocast: time-stable geocast for ad hoc networks," in *Proc. of VANET '05*, 2005.
- [4] R. Meier and V. Cahill, "STEAM: Event-Based Middleware for Wireless Ad-Hoc Networks," in *Proc. of DEBS '02*, 2002. [Online]. Available: citeseer.ist.psu.edu/meier02steam.html
- [5] S. Eichler, C. Schroth, T. Kosch, and M. Strassberger, "Strategies for context-adaptive message dissemination in vehicular ad hoc networks," in *Proc. of V2VCOM'06*, Jul. 2006.
- [6] C. Adler, R. Eigner, C. Schroth, and M. Strassberger, "Context-Adaptive Information Dissemination in VANETs Maximizing the Global Benefit," in *CSN*, 2006.
- [7] M. Caliskan, D. Graupner, and M. Mauve, "Decentralized discovery of free parking places," in *Proc. of VANET'06*, 2006.
- [8] T. Kosch, C. Schwingenschlgl, and L. Ai, "Information Dissemination in Multihop Inter-Vehicle Networks - Adapting the Ad-hoc On-demand Distance Vector Routing Protocol (AODV)," in *The 5th Int. Conf. on Intelligent Transportation Systems*, 2002.
- [9] D. Sormani, G. Turconi, P. Costa, D. Frey, M. Migliavacca, and L. Mottola, "Towards Lightweight Information Dissemination in InterVehicular Networks," in *Proc. of VANET 2006*, 2006.
- [10] ETH Vehicular Traces, "<http://lst.inf.ethz.ch/ad-hoc/car-traces/>"
- [11] GMSF generator, "<http://gmsf.hypert.net/>"
- [12] A. Vahdat and D. Becker, "Epidemic Routing for Partially Connected Ad Hoc Networks," Duke University, Tech. Rep., 2000.
- [13] V. Bychkovsky, K. Chen, M. Goraczko, H. Hu, B. Hull, A. Miu, E. Shih, Y. Zhang, H. Balakrishnan, and S. Madden, "The CarTel mobile sensor computing system," in *SenSys '06*, 2006.
- [14] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. on Information Theory*, 2000.
- [15] OMNET++ Website, "<http://www.omnetpp.org/>"
- [16] M. F. for OMNeT++, "<http://mobility-fw.sourceforge.net/>"
- [17] I. Leontiadis and C. Mascolo, "Opportunistic Spatio-Temporal Dissemination System for Vehicular Networks," in *Proc. of MobiOpp'07*, 2007.
- [18] G. Korkmaz, E. Ekici, F. Ozguner, and U. Ozguner, "Urban multi-hop broadcast protocol for inter-vehicle communication systems," in *Proc. of VANET'04*, 2004.
- [19] B. Xu, A. Ouksel, and O. Wolfson, "Opportunistic resource exchange in inter-vehicle ad-hoc networks," in *Proceedings of the IEEE International Conference on Mobile Data Management (MDM 2004)*, 2004.
- [20] S. Dornbush and A. Joshi, "StreetSmart Traffic: Discovering and Disseminating Automobile Congestion Using VANET's," in *Proceedings of the 65th Vehicular Technology Conference*, Dublin, Ireland, April 2007.