Opportunistic Spatio-Temporal Dissemination System for Vehicular Networks

Ilias Leontiadis
Department of Computer Science
University College London
Gower Street, London, WC1E 6BT
United Kingdom
I.Leontiadis@cs.ucl.ac.uk

Cecilia Mascolo
Department of Computer Science
University College London
Gower Street, London, WC1E 6BT
United Kingdom
C.Mascolo@cs.ucl.ac.uk

ABSTRACT

Opportunistic dissemination protocols have potentially applications in the domain of vehicular networking, ranging from advertising to emergency/traffic/parking information spreading: one of the characteristics of vehicular networks is that they are often partitioned due to lack of continuity in connectivity among cars or limited coverage of infostations in remote areas. Most available opportunistic, or delay tolerant, networking protocols, however, fail to take into account the peculiarities of vehicular networks.

This paper introduces a novel opportunistic event dissemination protocol for vehicular networks. The protocol takes into account the characteristics of these networks in order to dispatch the publications to the subscribers. Furthermore, it uses opportunistic cache and replay mechanisms to deliver the notifications to new subscribers in the area throughout the publication interval. We evaluate our approach through simulation using realistic vehicular traces. We compare our algorithm with a standard epidemic protocol, which offer the best alternative in terms of message delivery, by measuring overhead and delivery over a number of scenarios.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*; D.2.8 [**Computer-Communication Networks**]: Network Protocols—*Routing protocols*

General Terms

Algorithms

Keywords

Opportunistic Routing, Vehicular Networks, Information Dissemination, Publish/Subscribe

MobiOpp'07, June 11, 2007, San Juan, Puerto Rico, USA.

1. INTRODUCTION

More and more vehicles are now equipped with a satellite navigation system, allowing their geographical location to be readily available. Furthermore, it is expected that in the near future vehicles will be equipped with wireless interfaces, enabling them to form mobile ad-hoc networks on the fly and connect to fixed infostations while passing by [7]. *Vehicular networks* are hybrid mobile ad-hoc networks where infostations and vehicles are present. Infostations are fixed access points, potentially connected to the backbone (e.g. Internet). They may act as collection and dissemination points, from where information from the backbone network can flow towards the vehicles and information from remote vehicles can reach the backbone. Vehicles inter-network with each others, disseminating messages further in remote areas where there is poor or no coverage.

There is a growing number of applications for vehicular networks that require to disseminate information around specific geographical areas. For example, a driver might require to receive available parking spaces around his current location (in real-time). Or, to be able to monitor traffic conditions along his/her suggested route: In case of a traffic jam the navigation system may automatically calculate an alternative route to the driver's destination. Other examples include receiving nearby accident warnings, advertisements (e.g., fuel prices from local stations), vehicle synchronisation (e.g., merging lanes) etc.

In this paper we propose an approach which allows for message dissemination to all interested receivers in an area. Our mechanism keeps the message *alive* in the area for a specified period of time. In the system we envisage, drivers can register to receive (i.e., subscribe) information of their interest; however, a large number of the subscriptions will be automatically generated by the navigation system (NS). The navigation system can use the destination of the driver, the suggested route and map information to decide if a vehicle is interested on receiving or not a specific notification. For example, the NS can subscribe to receive notifications about road works that affect the calculated route, subscribe to receive fuel prices from nearby fuel stations when the vehicle is running out of fuel, or to receive free parking notifications concerning a specific destination. The publisher can either be in the specified area already, or route the message to the interested area using some other approaches such as [10]. For simplicity, in this paper we assume that the publisher wishes to disseminate a message in the area around itself, but extensions are possible. Centralised solutions could of course be used but have many drawbacks:

- Centralised data may be outdated and the response time may not meet the real-time requirements.
- Current centralised communication solutions (GSM, WiMAX) may not be able to handle the burden of real-time monitoring of hundreds of thousands of vehicles.
- Infrastructure could be quite expensive, especially if the area to be covered is large.
- Infrastructure may not be available, especially in remote areas.

The use of an opportunistic networking approach, possibly exploiting geographical information might be preferable. This solution may provide real-time local information to the driver, it may be free of charge and would not require full coverage by infostations.

Recent research on vehicular information dissemination protocols has addressed some of these issues; in Abiding Geocast [11] the authors introduce a time-stable Geocast to a group of nodes in an 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 subscribers and topics in order to filter and route the notification to the interested nodes (it is an opportunistic Publish/Subscribe system and not a Geocast protocol). Furthermore, we exploit the unique characteristics of vehicular networks like mobility patterns, road topology, map availability and suggested routes to disseminate the notifications more efficiently. Finally, our evaluation shows that periodic broadcasts in a vehicular network are more efficient than epidemic spreading. In [2] the authors use periodic broadcasts for free parking information dissemination and aggregation in order to inform cars about the availability of parking positions in an area. In context-adaptive information dissemination [1], the authors highlight the fact that a danger warning might not be relevant to all areas. Therefore, they use a relevance factor to control the dissemination of a message. This function is affected by the distance from the event and time since the event was published. Similar research has been also conducted by T. Kosch et al. [9] in order to improve the AODV routing protocol for vehicular networks. Finally, in [13], the authors show how vehicles moving in the opposite direction can be exploited as carriers to quickly disseminate information to the vehicles that follow. None of these works, however, consider subscriptions to filter and disseminate the notifications to interested nodes and they do not allow for information to be maintained in an area (i.e., "to stick to an area") for a specified time. Furthermore, they do not exploit the unique characteristics of vehicular networks for the dissemination.

The paper is organised as follows: In Section 2 we briefly describe the publish/subscribe paradigm that we employ to efficiently disseminate the notification to the interested vehicles. Furthermore, we make some interesting observations about the distribution of subscribers in the notification area and we present our notification dissemination protocol. In Section 3 we present the results of our evaluation and show the performance under different conditions and by using realistic traffic traces generated by a multi-agent microscopic traffic simulator (MMTS) developed by K.Nagel at ETH, Zurich [14, 5]. The results show good performance in various settings in terms of overhead and message delivery, also with respect to epidemic dissemination [16], which we used as a benchmark. Section 4 contains conclusions and possible future work.

2. THE APPROACH

In this section we first outline the various steps of our algorithm and then we go into details of each of the steps. In brief, the algorithms can be described as follows:

1. Travellers and the vehicles' navigation system can subscribe to event types at any time (they can also unsubscribe).
2. A publisher emits an event (or notification), which, among others, contains the type of the event and the location of the event (e.g., parking spot). Furthermore, it indicates the area[1] in which it needs to be disseminated and the time validity of the event.
3. A number of event replicas are maintained in the area throughout the dissemination period in order to notify all the necessary subscribers. This number is tuned based on density information.
4. Event replicas are passed from vehicle to vehicle: the algorithm takes advantage of subscriptions and vehicular movement (which can be extracted from modern navigation systems) in order to make right guesses of where to disseminate the replicas.

2.1 Publish/Subscribe

As described in the outline above, the notification should be continuously disseminated around the publisher for a specific time interval. In order to clarify how application developers will interface with our protocol we briefly define the `notify()` primitive:

```
notify(message, topic, range, timeToLive)
```

where, `message` is the body of the message, `topic` is the type of the event, `range` is the range around the publisher that this notification should reach (i.e., the area) and `timeToLive`, which is the validity of the event. A driver or the vehicle's navigation system can subscribe to receive certain notifications using the subscribe primitive.

```
subscribe(topic)
```

The semantics of this primitive are simple: record an interest in the topic indicated. When a notification is received, the vehicle's matching engine determines if it is interested, not interested or already informed. As we are going to see, subscriptions are not just used to filter incoming notifications, they are also used to "route" the broadcasts towards the actual subscribers.

[1] For simplicity a radio range around itself, in this paper, but extensions to route the event towards an area are possible, for instance by using [10].

2.1.1 Topic

The topic should be able to accurately describe the context of the notification because when a notification is received the vehicle has to determine if it is interested in it or not, according to its subscriptions. There are a lot of existing solutions that allow this type of context definition and context matching [12, 18, 6], which we can take advantage of.

In our solution, the topic *contains the location(s) that the publication concerns (the point of interest)*. For example a parking spot notification contains the location of the spot. An accident warning contains the location of the warning. A fuel price advertisement contains the location of the fuel station. This context information is important for the driver or the navigation system in order to determine if it is interested for this notification or not. Additionally, the topic can also contain keywords that can be used to match a subscription for the needs of a specific application (e.g. Category = Road Warning, Advertisement etc.)

In order to organise this information, we selected a hierarchical structured topic definition. The publications and subscriptions can be matched at any level of the topic-trees. XML files can be used to describe such kinds of structure. The navigation system will combine this information with GPS and its maps to match a notification to a subscriptions.

2.2 Mobility Patterns and Distribution of Subscribers

Modern Navigation Systems (NS) typically consist of a Global Positioning System (GPS) device, maps, and the appropriate hardware and software. Their main function is to calculate a suggested route from the current position of the vehicle to the destination of the driver and provide turn-to-turn navigation assistance to the driver until the vehicle reaches the destination. The driver may select his/her destination and preferences (e.g., calculate fastest route, shortest route, avoid highways/tolls, etc), and the navigation system calculates a *suggested route* from the current position of the vehicle to the final destination.

In the vehicular scenarios that we have just described, the notification concerns only a specific geographic point. For example, the notification about the free parking spot concerns only vehicles that drive towards this point. Similarly, a traffic jam only affects vehicles that plan to drive through it. Fuel prices of a specific station concern only vehicles driving towards it (e.g., not the ones in the opposite lane). It is obvious that there are going to be very few subscribers on road segments that do not lead to the *point of interest*. In other words, *the subscribers are not evenly distributed inside the publication area*. This distribution is affected by the position of the point of interest, the mobility patterns of the vehicles and the road topology. To optimise communication, a mechanism to route the notification towards the areas that contain actual subscribers -and not any vehicle- is needed.

2.3 Notification Dissemination

Not every vehicle is broadcasting the notification. Instead, a (small) number of replicas N_{rep} are created. These replicas of the message will be the notification "broadcast points". The vehicles carrying these replicas will become *mobile info-stations* for the notification and will periodically broadcast it. Our aim is to create and distribute an adequate number of replicas so that their periodic broadcasts cover the areas where subscribers are likely to be. As we will show (Section 2.3.2), the number of replicas is dynamically adjusted, but, in general, this is much lower than the number of subscribers. Vehicles carrying a replica and moving outside the notification area have to transfer the replica to a vehicle that is driving back inside. To keep informing all the subscribers in the area (including the new ones joining any time during the notification period), *every replica is broadcast every F_{int} seconds. Every subscriber within the broadcast radius r of a replica is then notified.*

Simultaneously to the broadcast, the replica may be transferred to a new vehicle: the replica holder designates another node as new replica holder and deletes his copy after the broadcast. We call this operation *replica forwarding*. The motivation of this strategy is simple: we cannot use just the mobility patterns of the vehicles to route and distribute the notification to the subscribers. For example, let us imagine that the publication area contains only a one-way segment of a highway. If the replicas were always kept in the same nodes (that periodically re-broadcast them), then all the replicas would concentrate in roughly the same point: the boundary of the publication area where the vehicles exit the highway segment. Therefore, it is essential to forward the replicas from vehicle to vehicle in order to keep them as evenly distributed as possible and in areas that contain subscribers.

There are two challenges in this. Firstly, how to select the appropriate *number* of replicas N_{rep} to achieve high delivery ratio with the minimum message overhead, and secondly, how to *select the next replica carrier*. Notice that, for the sake if simplicity, we decided to keep the forwarding interval F_{int} as an application parameter and adapt the number of replicas to the networks conditions.

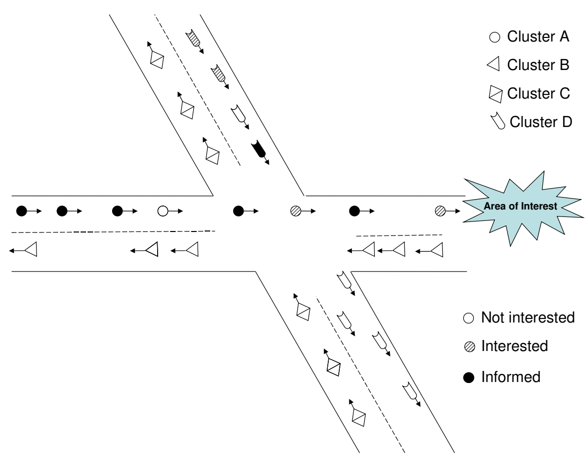

Figure 1: Clusters of cars and subscribers

2.3.1 Next Replica Carrier Selection

Vehicles *move in clusters towards a specific direction*. For example, on a two way road we can observe two main clusters: the cars moving in one direction and the cars moving in the opposite. Similarly, at a crossroad there can be two to eight clusters. In Figure 1, four clusters can be observed.

Just before broadcasting a notification, the replica owner broadcasts a message containing the subscription topic (but not the message itself). All one hop neighbours respond with their *subscription status, position and direction.* [2] The *subscription status* indicates whether the neighbour is interested in this notification. A neighbour responds with:

- `Informed:` when it has received that notification before.
- `Interested:` when there is a subscription that covers the topic. The navigation system and its subscription topics are consulted to decide if the vehicle is interested or not in the specific notification (e.g. match if the point of interest, the current suggested route and the topic match)
- `Not Interested:` otherwise.

After collecting the replies, the current replica owner uses the *direction* information to cluster the neighbour vehicles. Although the map can be consulted for more accurate results, a simple angle threshold algorithm is also adequate: In our simulations we consider that a node belongs to a cluster if its direction is deviating a maximum of ± 5 degrees from the cluster average (we add the node to the cluster and update the cluster average. If there is no such cluster we create one and add the host). This approach is efficient if you consider that cars move on roads towards specific directions.

To *select the next carrier*, the replica owner *finds the cluster that has the most uninformed subscribers* (i.e., subscribers that responded with *Interested*). If there is no such cluster with uninformed subscribers, the replica owner finds the cluster with the most subscribers. It then selects as the next replica carrier *a random vehicle in a cluster that is moving towards the opposite direction.*

The motivation for this is that a replica is needed in the area where the uninformed interested subscribers are coming from. The ideal would be to select a vehicle that is going towards that area. For example, Figure 1 illustrates two clusters that contain subscribers. These clusters contain vehicles that are driving toward the area of interest. However, most of the vehicles of cluster A are informed (by another replica). Cluster D has the most uninformed subscribers. Therefore a replica needs to be sent towards the area where these cars are coming from. Hence we select a vehicle from cluster C as the next carrier.

This mechanism encourages the replicas to be routed towards areas with uninformed subscribers. Furthermore, if two or more replica holders are broadcasting in the same road section, one of these will eventually move somewhere else because there are not going to be enough uninformed subscribers to keep it there.

However, the next carrier is not always selected that way. When the replica owner is close to the area boundaries a node that is moving towards the centre is selected so as to keep the replica inside the area.

2.3.2 Number of replicas

In order to achieve a high delivery ratio with a minimum message overhead, an appropriate number of replicas N_{rep} needs to be selected. Our aim is to delete replicas that are not useful (e.g., located in areas where there are no uninformed subscribers) and create more replicas in areas where there is a large number of uninformed subscribers. Furthermore, we do not want to create too many replicas because their broadcasts will overlap (i.e., the same notification will be delivered multiple times to the same subscribers).

As described in the previous section, before each broadcast, the replica owner polls the neighbours for their subscription status in order to select the next carrier. This information can be also used to estimate if the replica is needed in the area; If the replica meets uninformed subscribers then it is apparently needed in this area. If the replica does not meet any subscribers or all the subscribers that it meets have already been informed then it is not needed and should be deleted.

To calculate a more accurate estimation, every replica holder keeps the results from the last k polls (*number of non informed subscribers.* Then we use these statistics to adapt the number of replicas:

- If for the last k polls the replica met at least one uninformed subscriber then the replica is kept.
- If for the last k polls the replica met more than k uninformed subscribers then an additional replica is created: the replica owner does not delete its copy after forwarding the replica.
- If the replica owner did not find any uninformed subscriber then this replica is marked for deletion. In order to avoid deleting too many replicas simultaneously, we merge a replica with another one: a replica is actually deleted only if it detects a broadcast of another.

This method ensures that not too many replicas are created in the notification area. If many replicas are present, some of them will not find uninformed subscribers and they will be deleted. Also, if some replicas are in areas where there are no subscribers at all, they will also be deleted. In areas where there is a large flow of uninformed subscribers additional replicas will be created. Furthermore, notifications that concern a high number of vehicles will be broadcasted by more replicas whereas notification that concern very few vehicles will not create much overhead.

3. EVALUATION

In order to evaluate our approach, we used *OMNet++*[17, 15], a discrete event simulation environment and the *mobility framework plug-in* [4], which supports node mobility, dynamic connection management and a wireless channel model. We also used realistic vehicular traces in order to make this simulation as realistic as possible.

The aim of this evaluation is to validate how much the performance of our protocol benefit from the use of mechanisms for message replication to disseminate information only to relevant areas. In this sense, we evaluate the sensitivity to different scenarios, including different area size and density. We compare our protocol with an epidemic dissemination protocol, which offers very good message delivery at the price of a high overhead. We also compare our carrier selection choice with random choice.

[2] Note that the polling and the neighbour reply packet size is significantly smaller than the notification packet (that contains the message of the notification). More details about this in Section 3.

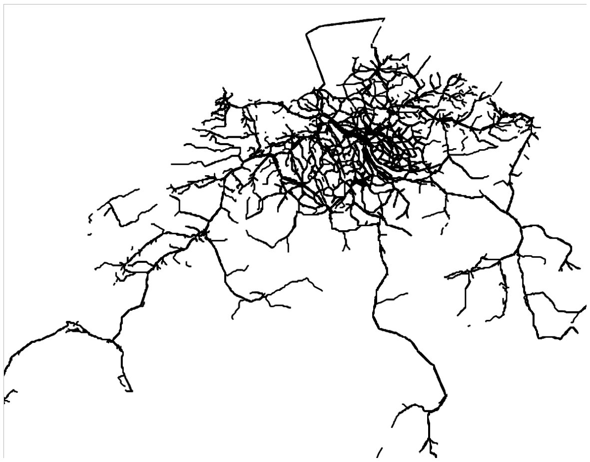

Figure 2: Map of the Vehicular Traces.

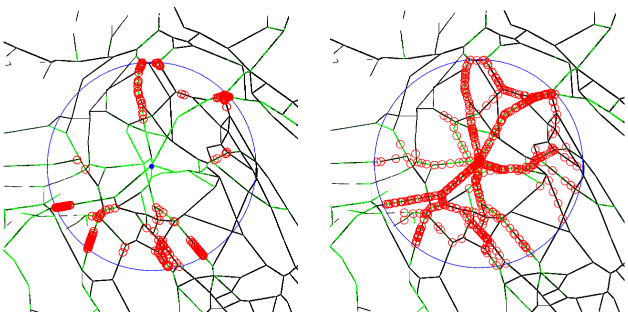

(a) Our carrier selection (b) Random carrier selection

Figure 3: Distribution of broadcasts. Green (or light grey) road segments contain high number of subscribers.

3.1 Vehicular Traces

In order to accurately evaluate our protocol in the context of vehicular networking, it would not make much sense to use any random mobility model [3] (random path, random waypoint, etc). Because no large scale vehicular traces exist, we have evaluated our approach by using traffic traces generated by a multi-agent microscopic traffic simulator (MMTS) developed by K.Nagel at ETH, Zurich [14, 5]. These traces contain mobility patterns of 260.000 vehicles over real road maps in the canton of Zurich within a period of 24 hours. Furthermore, they contain dense populated areas (the city of Zurich) and the surrounding highways, which enable us to run our simulations in different settings (Figure 2).

For our evaluation we decided to extract smaller areas from the 250km x 260km area of the traces. The reason is that simulation of 260.000 vehicles makes the simulation extremely slow. Furthermore, we wanted to extract many different scenarios with different characteristics (e.g., city scenarios, highway scenarios, etc). The areas that we selected are 20km x 20km. Figure 3 illustrates such a scenario.

3.2 Simulator settings

For our simulation we assume that 802.11b [8] wireless radio interface are used. The maximum possible communication range is 250m. All the broadcasts occur at the same channel frequency and the mobility framework determines which vehicles are able to receive the transmission by evaluating the signal to noise ratio. We used this radio interface because it is similar to the new 802.11p standard (Wireless Access for the Vehicular Environment or WAVE) and because it is already widely used both for simulation and in real life. The notification packet payload size (topic + message) is set to 1Kb, the polling message (topic) is 30bytes, the reply message is 10bytes and the message for merging replicas (F_{int}) is 4bytes. The mobility framework also adds the 802.11b broadcast headers to these messages. The forward interval is 30 seconds and the initial number of replicas are 50. The publication time to live is 3600 seconds (one hour).

The notification area radius is 7km (except at the experiments where we evaluate the radius). In the centre of the notification area we created a point of interest. Subscribers are all the vehicles with routes that will drive them less than 1km from the point of interest (i.e., will drive near the point of interest). In Figure 3, you can see the notification area and the point of interest. The green (or light grey) areas contain a high number of subscribers. In particular, in this scenario 25% of the vehicles are subscribers. The results that we present are averages of ten runs.

3.3 Results

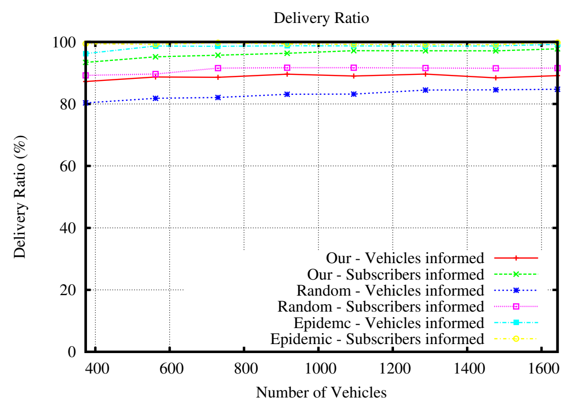

Figure 4: Delivery Ratio to subscribers only or to any vehicle for different densities.

The following experiments aim to evaluate how our notification spreading algorithm (described in Section 2) works in a variety of settings. To evaluate our carrier selection mechanism we compare it with random selection: the network is still polled in order to adapt the number of replicas but the replica carriers forward the message to a random neighbour instead of selecting a neighbour using the clustering technique. We also implemented epidemic message dissemination: when an infected (or informed) vehicle reaches another vehicle not met before, it sends a small poll message to check if the vehicle has been informed. If it has not, it sends the notification. We modified the epidemic protocol so as to constrain the dissemination to the publication area (vehicles outside do not infect others).

In Figure 3 we have also plotted (a random subset of) the

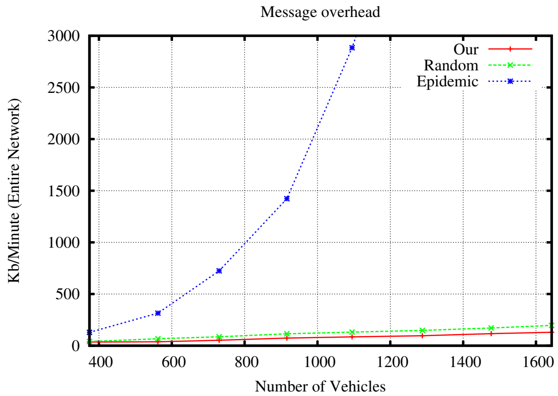

Figure 5: Message overhead for different densities.

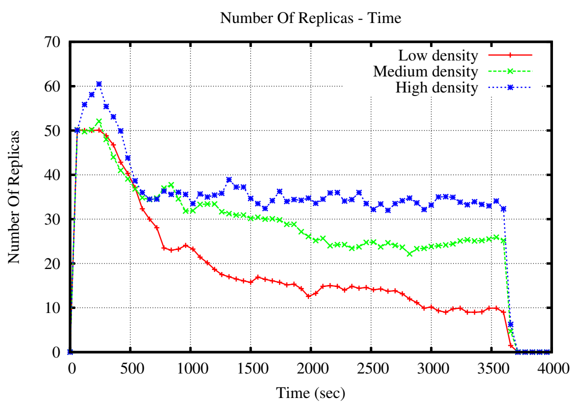

Figure 6: Number of replicas through time for different densities.

broadcasts that occur during the simulation on the map of the vehicular scenario that we used. As we notice, when random neighbour selection is used the broadcasts may occur anywhere inside the notification area. However, when we use our carrier selection algorithm, the broadcasts are routed towards the areas with the most uninformed subscribers. These areas are the entry areas of the subscribers to the system. By doing so a) the notification is delivered as soon as possible to the incoming vehicles and b) the notification is not delivered multiple times as the vehicle travels inside the notification area.

Figure 4 shows the delivery ratio of the notifications to the *subscribers only* and to *any vehicle* for varying vehicle density. When random carrier selection is used, the delivery ratio to the subscribers and to any host are almost the same. The justification is that the broadcasts are occurring in random places in the publication area. When we use our carrier selection mechanism we observe that the delivery ratio to any subscribers only is much higher. This indicates that *our mechanism efficiently routes the broadcasts toward the actual subscribers*. Epidemic message dissemination achieves the best delivery ratio both to subscribers and any vehicle. However, as we can observe in Figure 5, epidemic message dissemination introduces considerably higher overhead. This happens because epidemic dissemination requires that all the vehicles that are infected poll all the vehicles not met before. In our protocol, this number depends on the number of replicas and the forward interval only (not the number of vehicles inside). Furthermore, our protocol delivers the notification to more than one vehicle per broadcast. As we observe, in our algorithm the number of broadcasts increases linearly with the density due to the fact that more replicas have to be created to compensate the increased flow of uninformed subscribers. Furthermore, we notice that the random carrier selection protocol requires more replicas (and thus overhead) due to the fact that many replicas move in areas where there are no uninformed subscribers. In Figure 6 we observe how our algorithm adapts the number of replicas during the simulation for different densities. At the end of the notification area (second 3600) the replicas are deleted and the notification broadcasts stop.

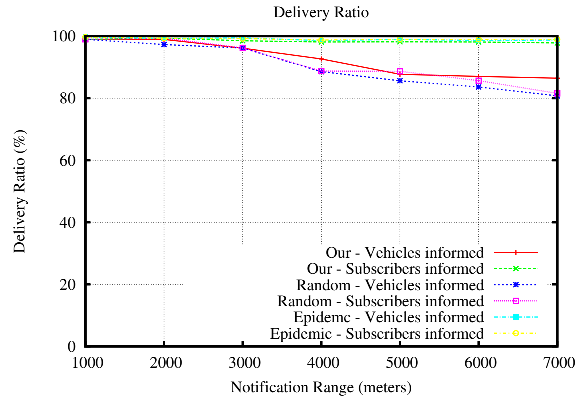

Figure 7: Delivery Ratio to subscribers only or to any vehicle for different notification range.

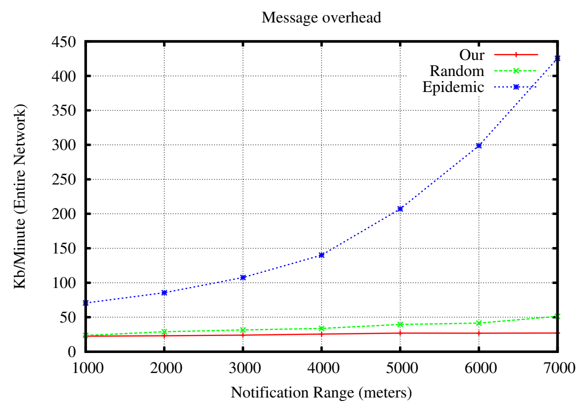

Figure 8: Message overhead for different notification range.

In Figure 7, we observe that for small notification areas (i.e., four times the radio range) the three algorithms have the same high performance because even a few replicas can keep informing the notification area. However, for larger notification areas, where the position and the number of the broadcasts matters, we observe that only our approach and epidemic message dissemination achieve acceptable performance. However, as we see in Figure 8, the message overhead of epidemic is again much higher due to the high number of polls. Furthermore, as we experimented with many locations, we noticed that the number of necessary replicas increases with the area size for the city scenarios but it remains almost constant for highway scenarios. This occurs because the number of required replicas is mainly affected by the number of locations containing uninformed subscribers; we do not need more replicas for the same road segment.

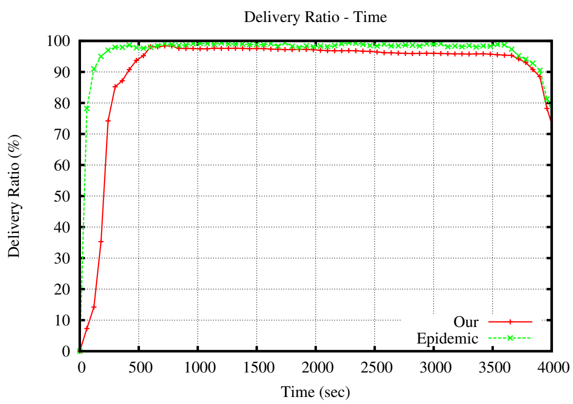

Figure 9: Delivery Ratio to subscribers through time.

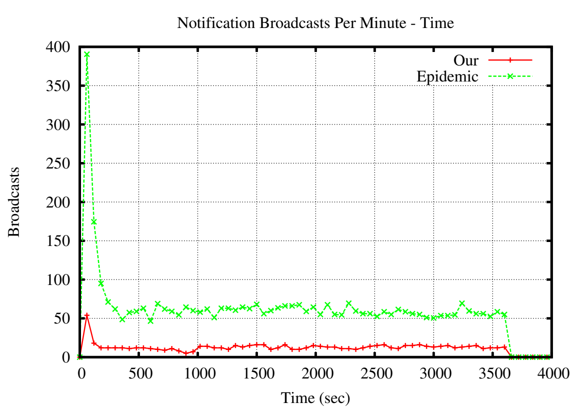

Figure 10: Number of broadcasts of the notification through time.

Figure 9 illustrates the delivery ratio throughout the simulation. As we observe, the epidemic protocol quickly starts

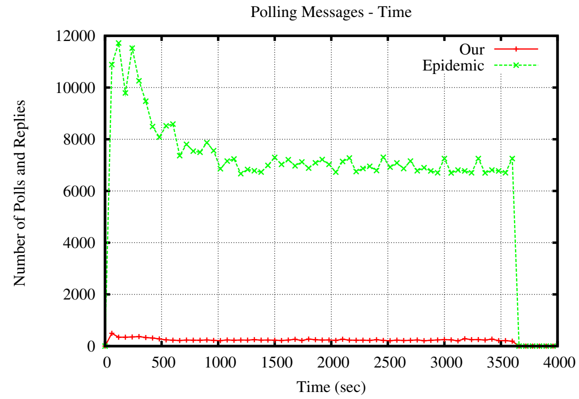

Figure 11: Number of polls and replies through time.

to infect all the vehicles in the area. However, in our mechanism the replicas require some time to be routed into the correct areas before we reach a high delivery ratio. Additionally, in Figure 10 we notice that the number of broadcasts in epidemic is much higher at the beginning of the notification period but the anti-entropy slows this down as more and more vehicles become infected. Some minutes after, epidemic still needs more broadcasts per minute to inform the new subscribers in the area. Moreover, in Figure 11 we also notice that the epidemic protocol polls a significantly higher number of vehicles for their notification status.

4. CONCLUSIONS

In this paper we have illustrated an opportunistic spatio-temporal dissemination system for vehicular networks:

- We introduced the publish/subscribe communication paradigms to opportunistically disseminate such information in vehicular networks inside a specified geographical location for the duration of a certain time interval.
- We took advantage of the information that can be extracted from the vehicle's navigation systems (location, map, destination of the driver etc) to generate subscriptions.
- We have shown that the mobility patterns of the vehicles and their subscriptions are actually linked (due to their geographical nature).
- We showed that we can take advantage of this observation to efficiently disseminate events only in areas where this is relevant (i.e., in areas where there are hosts that are interested in receiving the notification).
- We limit the overhead of the mechanism to the number of interested hosts by adapting the number of replicas that broadcast the notification using local observations.

We have evaluated our approach using realistic traffic traces generated by a traffic simulator. The simulation results indicate that our algorithm can efficiently deliver the notification in the area, taking advantage of the distribution of subscribers and to route the broadcasts towards them. We have also noticed that the number of broadcasts is mainly affected by the number of road segments that contain subscribers and the density of subscribers in the area. Finally, we have shown that our algorithm can achieve almost 100% of delivery ratio with less message overhead than epidemic due to the fact that it a) disseminates the notification mainly to subscribers b) it can inform multiple subscribers per broadcast and c) it has less polling overhead.

In the interest of keeping things simple, this paper has not considered optimisations of routing based topic aggregation. As future work, we would like to extend our evaluation to asses how existing topic aggregation techniques could improve performance further. Furthermore, we did not consider routing the notification to remote areas (for example sending traffic alerts from an infostation to be disseminated in a specific geographical region).

Acknowledgements: We would like to thank Paolo Costa, Vladimir Dyo, Liam McNamara, Mirco Musolesi, Daniele Quercia and David Rosenblum for comments on an earlier draft. We also acknowledge the support of the EPSRC through Project CREAM.

5. REFERENCES

[1] Christian Adler, Robert Eigner, Christoph Schroth, and Markus Strassberger. Context-Adaptive Information Dissemination in VANETs Maximizing the Global Benefit. In *Fifth IASTED International Conference on Communication Systems and Networks (CSN 2006)*, 2006.

[2] Murat Caliskan, Daniel Graupner, and Martin Mauve. Decentralized discovery of free parking places. In *VANET '06: Proceedings of the 3rd international workshop on Vehicular ad hoc networks*, pages 30–39, New York, NY, USA, 2006. ACM Press.

[3] T. Camp, J. Boleng, and V. Davies. A Survey of Mobility Models for Ad Hoc Network Research. *Wireless Communications and Mobile Computing (WCMC)*, 2(5):483–502, 2002.

[4] W. Drytkiewicz, S. Sroka, V. Handziski, A. Koepke, and H. Karl. A Mobility Framework for OMNeT++. 3rd International OMNeT++ Workshop, at Budapest University of Technology and Economics, Department of Telecommunications Budapest, Hungary, January 2003.

[5] ETH Zurich Vehicular Network Traces. http://lst.inf.ethz.ch/ad-hoc/car-traces/.

[6] Ludger Fiege, Felix C. Gärtner, Oliver Kasten, and Andreas Zeidler. Supporting Mobility in Content-Based Publish/Subscribe Middleware. In *Middleware*, pages 103–122, 2003.

[7] Ralf Guido Herrtwich. Communicating Vehicles - Communicating Roadways: New Approaches to Driver Information and Road Safety. In *Invited Talk at MobiCom*, 2005.

[8] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Higher-speed Physical Layer Extension in the 2.4 GHz Band. Technical Report IEEE Std 802.11b-1999 (R2003), IEEE, 2003.

[9] Timo Kosch, Christian Schwingenschlgl, and Li Ai. Information Dissemination in Multihop Inter-Vehicle Networks - Adapting the Ad-hoc On-demand Distance Vector Routing Protocol (AODV). In *The IEEE 5th International Conference on Intelligent Transportation Systems*, Singapore, 2002.

[10] Ilias Leontiadis and Cecilia Mascolo. GeOpps: Opportunistic Geographical Routing for Vehicular Networks. In *Proceedings of the IEEE Workshop on Autonomic and Opportunistic Communications. (Colocated with WOWMOM07)*, Helsinki, Finland, June 2007. IEEE Press.

[11] Christian Maihofer, Tim Leinmuller, and Elmar Schoch. Abiding geocast: time–stable geocast for ad hoc networks. In *VANET '05: Proceedings of the 2nd ACM international workshop on Vehicular ad hoc networks*, pages 20–29, New York, NY, USA, 2005. ACM Press.

[12] R. Meier and V. Cahill. STEAM: Event-Based Middleware for Wireless Ad-Hoc Networks. In *1st International Workshop on Distributed Event-Based Systems (DEBS '02), Vienna, Austria*, 2002.

[13] Tamer Nadeem and Pravin Shankarand Liviu Iftode. A comparative study of data dissemination models for vanets. In *3rd Annual International Conference on Mobile and Ubiquitous Systems (MOBIQUITOUS)*, 2006.

[14] Valery Naumov, Rainer Baumann, and Thomas Gross. An Evaluation of Inter-Vehicle Ad-Hoc Networks Based on Realistic Vehicular Traces. In *MobiHoc '06: Proceedings of the sSventh ACM International Symposium on Mobile Ad-Hoc Networking and Computing*, pages 108–119, New York, NY, USA, 2006. ACM Press.

[15] OMNET++ Community Website. http://www.omnetpp.org/.

[16] A. Vahdat and D. Becker. Epidemic Routing for Partially Connected Ad Hoc Networks. In *Technical Report CS-200006, Duke University*, 2000.

[17] András Varga. The OMNET++ Discrete Event Simulation System. *Proceedings of the European Simulation Multiconference (ESM'2001)*, 2001.

[18] Eiko Yoneki and Jean Bacon. An Adaptive Approach to Content-Based Subscription in Mobile Ad Hoc Networks. In *PERCOMW '04: Proceedings of the Second IEEE Annual Conference on Pervasive Computing and Communications Workshops*, page 92, Washington, DC, USA, 2004. IEEE Computer Society.