

When Cars Start Gossiping

Paolo Costa
Vrije Universiteit Amsterdam
The Netherlands
costa@cs.vu.nl

Daniela Gavidia
Vrije Universiteit Amsterdam
The Netherlands
dgavidia@cs.vu.nl

Boris Koldehofe
IPVS - Universität Stuttgart
Germany
Boris.Koldehofe@ipvs.uni-
stuttgart.de

Hugo Miranda
University of Lisbon
Portugal
hmiranda@di.fc.ul.pt

Mirco Musolesi
Dartmouth College
USA
musolesi@cs.dartmouth.edu

Oriana Riva
ETH Zürich
Switzerland
oriva@inf.ethz.ch

ABSTRACT

Vehicular ad hoc networks present challenging characteristics, such as very dynamic behavior and sparse connectivity, that need to be taken into account in designing adequate communication support. Gossip-based protocols have recently emerged as an effective approach to providing reliable and efficient communication in this domain. Nonetheless, despite the preliminary encouraging results, to the best of our knowledge, no previous work has systematically analyzed how gossip protocols are affected by the intrinsic characteristics of vehicular networks such as the very specific mobility patterns of vehicles, the relative abundance of memory and computational resources that vehicles offer, and the availability of geographical information through GPS receivers. In this paper, we aim at filling this gap by examining core requirements of vehicular network applications and analyzing the research challenges that gossip-based communication protocols need to address.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; H.3.3 [Information Search and Retrieval]: Retrieval models; H.3.4 [Systems and Software]: Distributed systems

Keywords

Vehicular Ad Hoc Networks, Gossip protocols

1. INTRODUCTION

Communication using gossip protocols is analogous to rumor spreading in a social scenario or epidemic spreading in a population. Essentially, in gossip protocols, each participating node receiving a message for the first time retransmits it

to a known subset of the members of the network until it is disseminated to all or to a given portion of the nodes.

Due to their inherently decentralized and robust nature, gossip algorithms have been applied to many domains of distributed systems, such as data replication [7], scalable information dissemination [9, 18], and mobile computing [20]. Therefore, gossip protocols represent a promising solution also for dynamic and unpredictable environments such as Vehicular Ad hoc NETWORKS (VANETs).

VANETs are Mobile Ad hoc NETWORKS (MANETs) where interacting nodes are computers embedded in cars. VANET applications aim at enhancing the driver's experience as well as providing alerts to several types of driving hazards. Typical application examples include providing customized directions to free parking spots, predicting upcoming traffic jams, or warning about icy roads or local fog patches. Like MANETs, VANETs present unstable configurations, unknown network delays, and high volatility. However, unlike typically resource-constrained devices of MANETs, computers embedded in cars can rely on the vehicle's power supply and may have plenty of memory storage available. They can also determine their absolute position by means of widely available GPS receivers.

Both academia and industry have shown an increasing interest in building vehicular networking systems, thus leading to a plethora of research projects focusing on devising appropriate communication models exploiting car-to-car short-range wireless connectivity [1–3]. Recently, gossiping protocols have also been considered in this context [8, 17, 26, 28]. However, as gossip protocols for wired networks cannot be directly applied to MANETs [12], gossip protocols proposed for MANETs [13, 15] cannot be directly applied to VANETs. In order to develop gossip algorithms for VANETs, it is necessary to understand the peculiar characteristics of such networks. The potentially unlimited power supply of cars is such that the number of exchanged messages only becomes constrained by the available bandwidth. Furthermore, vehicular computers can provide resources for the execution of complex operations like context management, data filtering or even cryptographic algorithms. Hence, gossip communication in VANETs goes beyond requirements such as resource-efficient data communication that have been often addressed in the context of MANETs.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MINEMA'08, April 1, 2008 Glasgow, Scotland
Copyright 2008 ACM 978-1-60558-122-4/08/04 ...\$5.00.

This paper outlines how gossip communication can be successfully accomplished in VANETs and identifies major issues that still need to be considered. The rest of the paper is organized as follows. Section 2 motivates the use of gossip in supporting VANET applications while Section 3 presents core research challenges for achieving efficient gossip communication in VANETs. The paper concludes in Section 4.

2. WHY CARS SHOULD GOSSIP ?

Vehicular networks are expected to have a large impact in many domains of industry and society. For instance, by now a considerable amount of work has already focused on enhancing the driving experience and increasing road safety. Furthermore, industry looks into ways to better coordinate and optimize logistic processes also by means of vehicular networks. To accomplish easy and scalable deployment of such applications it is essential to provide suitable robust communication abstractions which can accommodate the specific requirements of vehicular applications.

In order to get a better understanding of the requirements involved, we present three realistic VANET scenarios.

Alice's Foggy Ride. During a night ride, Alice suddenly encounters a heavy fog patch that severely compromises her visibility, thus increasing the risk of accident. Alice is in a hurry and cannot wait for the fog to vanish. Therefore, she activates her traveling support system. The system has a radar-like display that shows the location of other vehicles and hidden road traps signaled by drivers ahead of her. To limit the level of distraction, Alice programs the device to inform her about traps at most 500 meters ahead. She also contributes to improve road safety by creating and forwarding alerts to drivers behind her.

Bob's Weekly Dilemma. Every week-end, Bob drives about 300 km to return home from his summer cottage. When approaching his home town, he is always faced with a dilemma: either to follow the shortest path, possibly heading to a traffic congestion, or to take the longer route, where traffic jams are rare. This week-end, Bob decides to use his enhanced route planner to learn about upcoming traffic jams along his favorite shortest path and, since no traffic jam is predicted, he can arrive home earlier than usual.

Charles's Enhanced Driving Experience. Charles works as salesman for a big company and he spends most of his time travelling across the country in remote areas. However, even though he is not familiar with those places, thanks to his enhanced navigator, he can always be aware of all nearby facilities. For instance, at lunch time, the navigator retrieves up-to-date information regarding nearby restaurants, including opening times and prices. This kind of information is provided and constantly refreshed by restaurant owners and can be received within the city limits. Also, if his car is running low on fuel, the closest gas station will be displayed on his on-board map with detailed indications about prices and different varieties of fuel sold.

Despite their simplicity, these application scenarios summarize most of the distinguishing features and system design issues of VANETs. Foremost, these networks are inherently characterized by high dynamicity both at the network level (since hosts move, possibly very fast) and at the application level (since the amount of information and data gener-

ated by vehicular applications may be very large, for example, traffic information requires continuous updates). In addition, their scale largely outperforms the size of traditional mobile systems, as they can consist of millions (if not hundreds of millions) of hosts. Finally, connectivity cannot be always guaranteed at any place, anytime.

These peculiarities simply preclude the use of established protocols like those used to route packets on the Internet, in cellular and WiFi networks. Indeed, all these protocols make strong assumptions about the stability of the underlying topology and the knowledge available at each host, which cannot be replicated in the VANET domain. Classic routing protocols based on forwarding mechanisms of packets adopted in MANETs [25] are also not suitable, given the high dynamicity of the network and the fact that a connected path may not always exist between the sender and receiver(s). Furthermore, in this kind of networks disconnections are the norm rather than the exception and store-and-forward mechanisms should be introduced to support delay-tolerant communication [11].

Consequently, in recent years, researchers have moved their attention towards gossip protocols [8, 17, 26, 28] as a viable alternative to address VANET challenges. Typical properties which have contributed to the popularity of gossip in various domains of distributed systems [9] are the simplicity in implementing the protocol, its ability to scale to a large number of participants, and its high resilience to failures. Furthermore, the probabilistic nature of these protocols deals well with the inherent non-determinism of vehicular networks, in which topology changes are frequent.

The underlying idea behind this family of algorithms is that each node has to periodically communicate its knowledge about the system "state" to a random subset of other nodes of the network. All gossip-based algorithms operate roughly in the same way: each node selects one or multiple nodes at random and sends them a message. The receiving nodes update their state and repeat the procedure. The state update and propagation mechanisms are the characterizing features of the different protocols that have been proposed in the literature [10]. Repeated execution of the protocol results in redundant messages which increase the probability of successfully propagating the information through the system.

Gossip protocols do not rely on any specific underlying topology and, hence, they can quickly adapt to ever-changing conditions. Notably, these protocols can even benefit from network dynamicity because vehicle mobility can be exploited to propagate and deliver messages in remote areas also in the presence of sparse networks. For instance, during his daily travels Charles can carry information from one area to another, thus effectively ensuring proper delivery even if no connected path exists between the two areas.

Finally, gossip protocols naturally support data aggregation and have been successfully used to compute network wide properties, e.g., the number of nodes with a specific property [21]. This is of paramount importance in VANETs given the vast amount of data available. For instance, in Bob's case, disseminating the location of every single car would prove extremely inefficient. Instead, traffic jams can be more efficiently detected using average measurements of cars' speed and density over multiple cars. Also in Alice's scenario, relying on information provided from a single source could not provide the required accuracy and confidence in the phenomena she tries to observe. Instead, she could definitely

benefit from aggregated information provided by cars that have observed the same phenomena.

Unfortunately, despite their suitability, the peculiarities of VANETs differ from the scenarios in which gossip protocols have traditionally been applied and impose an accurate redesign of these protocols. In the next section, we will illustrate the key challenges and discuss some possible solutions.

3. HOW CARS SHOULD GOSSIP ?

In this section, we discuss VANET-specific challenges that need to be addressed to devise efficient gossip protocols in such networks.

Limited Connectivity. A fundamental requirement of traditional gossip protocols is the ability to randomly select a node in the network. This randomness yields high convergence speed and robustness, since all nodes evenly take part in the process. Conversely, in VANETs, each node can communicate only with those nodes falling within its broadcast range, thus potentially affecting the overall dissemination. Therefore, robustness may not be guaranteed with traditional gossip mechanisms. For instance, in Bob's scenario, traffic jam alerts need to be spread across a vast area and local infection is not sufficient. On the other hand, hybrid approaches combining car-to-car and infrastructure-based communication (e.g., hot spots at gas stations) can be a solution as they allow for long-range communication and, at the same time, coverage in isolated areas.

Mobility Patterns. In VANETs, cars mobility is influenced by the territory geography and by the road infrastructure. Other constraints are related to the variable traffic density in different times of the day (i.e., rush hours) and to the presence of points of attraction, such as office buildings, malls, or schools. In other words, social aspects influence the mobility patterns. These mobility patterns lead to very peculiar dynamic topologies characterized by well-defined network structures, potentially affecting the performance of the dissemination process. Nonetheless, we believe that it will be possible to exploit recent results in complex network theory [4, 24] to improve the robustness and the performance of gossip protocols in terms of number of exchanged packets. It is difficult, however, to estimate such topologies in general, given the scarcity of wireless traces currently available.

On the other hand, cars mobility offers new opportunities to protocol designers because cars movement can be, at least to some extent, predicted using statistical traffic models (also related to social aspects) or logs provided by on-board navigation systems. Researchers have already exploited this kind of knowledge to drive message propagation [23]. For example, in Bob's scenario, cars moving towards highway on-ramps could be selected as good infective agents. Thus far, however, it is not clear how this *deterministic* selection deals with the *probabilistic* selection of traditional gossip protocols.

Geographical Information. We can safely assume that in the near future, every vehicle will be equipped with an on-board GPS system, which means timely information about current position and speed of the vehicle. This information can be exploited to develop more refined ways to scope the epidemic dissemination only within given areas [26]. In case of Alice's fog, for instance, notifications should be scoped within a given area to avoid overwhelming all drivers with useless information and to reduce communication overhead.

Opportunistic Routing. Traditional MANET routing protocols assume that the underlying network graph is always connected, i.e., given an arbitrary pair of nodes, it is always possible to find a multi-hop path between them. However, recently, researchers have started considering intermittently connected MANETs. These can be considered as an example of Delay Tolerant Networks (DTNs). The distinguishing feature of these networks is that connectivity is not guaranteed anytime. This is highly likely to occur in rural environments (like those ones visited by Charles) where it is unrealistic to assume a high density of cars or infostations. Yet, message delivery is still possible by exploiting store-and-forward mechanisms such as those employed by epidemic protocols. Nodes can temporarily store packets and carry them across different areas, thus distributing them to nodes that were not previously reachable. A major problem of existing DTN protocols is their scalability as most of them rely on oracle-based [14] or table-based [19] routing protocols. Smart epidemic protocols represent a promising solution to be considered, as they are stateless [22] and, hence, more scalable.

Persistency. In many scenarios, vehicular applications may benefit from the use of persistent notifications, i.e., messages need to be maintained in the system for a certain interval of time to ensure that all interested vehicles will be eventually notified. A typical example is represented by Bob's scenario where road work alerts (e.g., "Highway A51 will be closed from 8 to 12 due to road maintenance") usually require to last for a given amount of time. In this case, pure (epidemic) dissemination protocols are not sufficient because vehicles arriving later in the area should be notified too, even in the presence of a very sparse network. Up to a certain extent, fixed infostations could support a solution by periodically reinitiating the broadcast of a message. Unfortunately, this approach is not suitable within rural areas where it is likely that such infrastructure is unavailable. Specific techniques must be consequently devised to ensure that messages remain in the area of interest, e.g., by continuously passing from a vehicle leaving the area to one entering it. Epidemic-style solutions have been proposed [17] because they are a natural way of keeping information by means of an *endemic* state of the infection.

Communication paradigms. Gossip protocols appear as a viable solution to address large-scale robust dissemination in VANETs. Nevertheless, some applications (e.g., Charles' navigation system) require more sophisticated paradigms, such as publish-subscribe and query-advertisement mechanisms. Nevertheless, in the field of peer-to-peer networks, also characterized by high variability and churn, several researchers have already devised gossip-based protocols to support these advanced paradigms [6, 27]. It is not straightforward, however, to apply them to VANETs, as node mobility seriously affects the delivery process. A possible approach could be to exploit again the aforementioned route planning in order to allow the system to know with a good degree of accuracy the next position of the vehicle. If we could rely on some fixed infrastructure (e.g., hot-spots in highways), a feasible solution is to store replies in a specific location, where the vehicle is supposed to pass nearby. Alternatively, if the infrastructure is not available, a different approach might be to deliver replies on vehicles which are approaching the initiator vehicles in the opposite direction. A key assumption of this kind of systems is the cooperation among the devices to

support the routing process [16].

Data management. Given the large variety of concurrent tasks that VANETs will need to support, it is desirable to aggregate data to optimize the global resource utilization and achieve improved performance and scalability. Notifications with overlapping requirements in terms of data or regions of interests (e.g., Bob's traffic alerts) should be merged. This process implies a phase in which new merged notifications are computed and distributed to the areas of interest based on their initial requirements. In-network data aggregation has been largely investigated in the context of static sensor networks [5]. Compared to sensor networks, the mobility of VANETs highly complicates the aggregation process. For example, it is necessary to avoid that a particular piece of information is aggregated multiple times at different nodes, thus negatively affecting the quality of the result.

Ensuring security and privacy is fundamental for the success of VANETs. Gossip algorithms are particularly vulnerable, given that vehicles are expected to periodically broadcast to untrusted users, leaking data like the intended destination, current location or user identity. This data may be combined to infer and expose drivers private information, to impersonate them or to disseminate bogus data. In VANETs, however, we expect the vehicles to have sufficient computer power to perform complex cryptographic operations, thus partly alleviating the problem.

4. SUMMARY

In this paper, we have presented the core research challenges related to the application of gossip-based techniques to support communication in VANETs. Stemming from an analysis of the specific requirements of vehicular applications and of the peculiarities of vehicular networks, we have outlined a possible research agenda where key research issues include limited connectivity, mobility patterns, opportunistic routing, and data management.

Acknowledgments

The work described in this paper was partially supported by the European Science Foundation (ESF) under the MINEMA project.

5. REFERENCES

- [1] CarTel Project. <http://cartel.csail.mit.edu/>.
- [2] FleetNet Project. <http://www.ccrle.nec.de/Projects/fleetnet.htm>.
- [3] Traffic View Project. <http://discolab.rutgers.edu/traffic>.
- [4] R. Albert and A.-L. Barabasi. Statistical mechanics of complex networks. *Review of Modern Physics*, 74, 2002.
- [5] D. E. Bhaskar Krishnamachari and S. Wicker. The Impact of Data Aggregation in Wireless Sensor Networks. In *Proceedings of DEBS'02*, July 2002.
- [6] P. Costa and G. P. Picco. Semi-probabilistic Content-based Publish-subscribe. In *Proceedings of the 25th International Conference on Distributed Computing Systems (ICDCS'05)*, June 2005.
- [7] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic Algorithms for Replicated Database Maintenance. In *Proceedings of PODC'87*, 1987.
- [8] S. Dornbush and A. Joshi. StreetSmart Traffic: Discovering and Disseminating Automobile Congestion Using VANETs. In *Proceedings of the 65th Vehicular Technology Conference*, April 2007.
- [9] P. T. Eugster, R. Guerraoui, S. B. Handurukande, A.-M. Kermarrec, and P. Kouznetsov. Lightweight probabilistic broadcast. *ACM Transactions on Computer Systems*, 21(4):341–374, 2003.
- [10] P. T. Eugster, R. Guerraoui, A.-M. Kermarrec, and L. Massoulié. From Epidemics to Distributed Computing. *IEEE Computer*, 37(5):60–67, 2004.
- [11] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proc. of SIGCOMM'03*, 2003.
- [12] R. Friedman, D. Gavidia, L. Rodrigues, A. Viana, and S. Voulgaris. Gossiping on MANETs: The Beauty and the Beast. *ACM Operating Systems Review*, 41(5), 2007.
- [13] Z. J. Haas, J. Y. Halpern, and L. Li. Gossip-based ad hoc routing. *IEEE/ACM Transactions on Networking*, 14(3):479–491, 2006.
- [14] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proceedings of SIGCOMM'04*, August 2004.
- [15] D. Kempe, J. Kleinberg, and A. Demers. Spatial gossip and resource location protocols. In *Proceedings of STOC'01*, pages 163–172, 2001.
- [16] S.-B. Lee, G. Pan, J.-S. Park, M. Gerla, and S. Lu. Secure incentives for commercial ad dissemination in vehicular networks. In *Proceedings of MobiHoc'07*, pages 150–159, New York, NY, USA, 2007. ACM.
- [17] I. Leontiadis and C. Mascolo. Opportunistic Spatio-Temporal Dissemination System for Vehicular Networks. In *Proceedings of MobiOpp'07*, June 2007.
- [18] M.-J. Lin and K. Marzullo. Directional gossip: Gossip in a wide area network. In *European Dependable Computing Conference*, pages 364–379, 1999.
- [19] A. Lindgren, A. Doria, and O. Schelen. Probabilistic Routing in Intermittently Connected Networks. *Mobile Computing and Communications Review*, 7(3), July 2003.
- [20] J. Luo, P. T. Eugster, and J.-P. Hubaux. Route driven gossip: Probabilistic reliable multicast in ad hoc networks. In *Proceedings of INFOCOM'03*, 2003.
- [21] L. Massoulié, E. L. Merrer, A.-M. Kermarrec, and A. Ganesh. Peer counting and sampling in overlay networks: random walk methods. In *Proceedings of the 25th ACM symposium on Principles of distributed computing (PODC'06)*, pages 123–132. ACM, 2006.
- [22] M. Musolesi and C. Mascolo. Controlled Epidemic-style Dissemination Middleware for Mobile Ad Hoc Networks. In *Proc. of Mobiquitous'06*, 2006.
- [23] T. Nadeem, P. Shankar, and L. Iftode. A Comparative Study of Data Dissemination Models for VANET. In *Proceedings of Mobiquitous'06*, pages 1–10, July 2006.
- [24] R. Pastor-Satorras and A. Vespignani. Epidemic Dynamics and Endemic States in Complex Networks. *Physical Review E*, 63(6), 2001.
- [25] C. E. Perkins. *Ad Hoc Networking*. Addison-Wesley, 2001.
- [26] D. Sormani, G. Turconi, P. Costa, D. Frey, M. Migliavacca, and L. Mottola. Towards Lightweight Information Dissemination in Inter-vehicular Networks. In *Proceedings of the 3rd Int. Workshop on Vehicular ad hoc networks (VANET'06)*, 2006.
- [27] S. Voulgaris, E. Riviere, A.-M. Kermarrec, and M. van Steen. SUB-2-SUB: Self-Organizing Content-Based Publish and Subscribe for Dynamic and Large Scale Collaborative Networks. In *Proceedings of WONS'06*, Les Menuires, France, January 2006.
- [28] B. Xu, A. Ouksel, and O. Wolfson. Opportunistic resource exchange in inter-vehicle ad-hoc networks. In *Proceedings of MDM'04*, pages 4–12, 2004.