

Virtual Access Points for Vehicular Networks

Nikolaos Frangiadakis,
Dept of Computer Science,
University of Maryland,
College Park, MD 20742, USA
+1 301 405 2687
ntg@cs.umd.edu

Daniel Câmara
Institut Eurecom
2229, route des Crêtes-
BP 193 06904
Sophia-Antipolis, France
+33 4 93 00 82 14
daniel.camara@eurecom.fr

Fethi Filali
Institut Eurecom
2229, route des Crêtes-
BP 193 06904
Sophia-Antipolis, France
+33 4 93 00 81 34
fethi.filali@eurecom.fr

Antonio A. F. Loureiro
Department of Computer Science
Federal University of Minas Gerais
Caixa Postal 702 30123-970
Belo Horizonte, Minas Gerais
+55 31 3499 5589
loureiro@dcc.ufmg.br

Nick Roussopoulos
Dept of Computer Science,
University of Maryland,
College Park, MD 20742, USA
+1 301 405 2687
nick@cs.umd.edu

ABSTRACT

This paper introduces the concept of Virtual Access Points (VAPs) for wireless Vehicular Ad-hoc Networks (VANETS). This new technique allows data dissemination among vehicles, thus extending the reach of roadside access points to uncovered road areas. Each vehicle that receives a message from an Access Point (AP) stores this message and rebroadcasts it into non covered areas. This extends the network coverage for non time critical messages. The VAP role is transparent to the connected nodes, and designed to avoid interference since each operates on a bounded region outside any AP. The experiments show the presented mechanism of store and forward at specific positions present a gain, in term of all the evaluated parameters.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Wireless communication

General Terms

Algorithms, Reliability

Keywords

VANET, Virtual Access Points, network coverage, caching

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, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Mobilware'08, February 12-15, 2008, Innsbruck, Austria.
Copyright © 2008 ACM 978-1-59593-984-5/08/02...\$5.00.

1. INTRODUCTION

Due to recent technological advances and the automobile industry's ever increasing efforts for security wireless communication between moving vehicles is currently a very active research area. Furthermore, in the future, we will be surrounded by ubiquitous wireless networks, and equipped with devices that will strive to remain perpetually connected as users move. A major challenge is the behavior of devices mounted on vehicles in areas where available Road Side equipments or Access Points are either sparse or not present. This paper presents a simple, yet powerful, technique to extend coverage to nodes outside hot spots areas.

DSRC is emerging as the standard for Vehicle to Vehicle and Electronic toll collection and thus, future vehicles are expected to be equipped with it. It is reasonable to assume that although safety applications will remain a priority, less critical applications will employ the unused bandwidth. Simultaneously, WiMax is considered as a new force that could act in parallel to help data dissemination. However, no matter the protocol used to forward traffic, the development of a new technique that bridges areas without direct access to any roadside equipment is of great importance.

To handle this issue, many solutions have been proposed, ranging from the use of more powerful access points, satellite coverage, to mobile APs [8]. Most approaches however, present two problems, namely scalability and cost. Firstly, equipping every road with access points is expensive. Secondly, satellite coverage, although important, does not scale for massive uplink/downlink communication. Our approach, based on a caching solution, provides a simple and inexpensive way to extend network coverage to uncovered areas.

A key aspect of our approach is how nodes cooperate, in a coordinated way, to increase the network coverage area. The mobile nodes cache messages originating from the APs, and are able to act as a Virtual Access Points (VAP) to other nodes in specific non covered areas. Thus, the nodes, collaboratively, help to forward packets to areas that had no traffic previously.

The rest of the paper is organized as follows; Section 2 discusses related work, whereas Section 3 formulates the proposed solution and outlines our analysis. Section 4 presents the experiments and results, and, finally, concluding remarks and future research are given in Section 5.

2. RELATED WORK

Data dissemination is a key issue in any network, and so it has been well studied. It can be push-based when the data are repeatedly broadcasted in a broadcast cycle, or pull based when data is explicitly requested by the users, or hybrid using mixed push-pull approach. Push or push-pull based data delivery is very attractive in environments like the ones considered and has been studied extensively in various environments [11,12,13]. The effect of mobility is also discussed in more recent papers [14,15,16].

The area of vehicular communication is very active. Various techniques and target applications have been studied. Most of them target the problem of improving safety through the use of wireless networks. For this reason, often the proposed techniques prioritize critical messages and allow non critical applications to use overhead bandwidth. This work is adapted to this context being able to use just the overhead bandwidth for delivering messages. This work focuses on push based applications that require the dissemination of either large objects or data streams, often requested by more than one user. The network architecture followed here is based in the basic multiple Infostation model introduced by Goodman et al. in [4]. However, we will refer to the Infostations as Access Points (APs).

The main objective of our work is to provide access to areas out of the coverage of APs. A very interesting work with similar objectives is the SPAWN system, introduced by Gerla et al. in [1],[2]. In these, the authors discuss how vehicles should interact to accommodate swarming protocols, such as BitTorrent traffic. In SPAWN, the nodes passing through APs collect data that they subsequently exchange among nearby nodes. As opposed to our system, SPAWN focuses on a restricted application that generates great volumes of traffic. Nodes are required to carry possibly useless to them traffic and the BitTorrent protocol is bandwidth intensive. Moreover, the number of retransmissions of a message in a vehicular network is estimated to be approximately 3 and so our gain from using the swarming protocol in this environment is non-optimal.

The Data Mule project [3] and the Message Ferrying scheme [17], designed for sensor networks, propose the use of mobile nodes to collect data from the sensors, buffer it, and deliver the collected data to a sink. As opposed to these works, we consider the problem not of retrieving data from the nodes, but of disseminating it to them. The MULEs (Mobile Ubiquitous LAN Extensions) and ferries utilize nodes navigating through the sensor network to collect data in “mobile caches”. According to the Data Mule project, all the nodes are fixed and only the cache

is mobile. On the contrast, in our scenario all nodes are mobile but we cannot affect their trajectories. Message Ferrying also considers mobile nodes but in that approach, as well as in [6] and [7], the nodes are required to follow specific paths and even move in order to help message delivering. The work presented in [7] proposes a multicast protocol for the highway environment where information dissemination through message flooding for VANET environments is proposed. Our proposal advocates that using of a more systematic approach for data dissemination, which is more bandwidth efficient. Zebrant [18] uses Zebras to carry historical data to the sink. In this project latency is not important. However, the most important differences are that the network is extremely sparse, the nodes are not moving in roads, their speed is limited and when two nodes meet, we consider there is enough time to exchange larger volumes of information.

Chen et al. [5] study network delay as a function of the number of cars and their velocity. The authors note that node mobility on highways can improve end-to-end transmission delay when messages were relayed. Furthermore, that low density networks may experience higher delays. These results are directly related to our work. VAPs locations should be selected so that information is not too widely spread and messages of time sensitive applications can reach their destinations in time.

An appealing solution to the problem, but more expensive and probably hard to implement, is the system proposed by Gavrilovich in [8]. In this work, to compensate for the velocity of the cars, the authors propose the creation of a chain of mobile APs on the center of the highway. Their role is to increase coverage and compensate for the high velocity of cars on the highway.

3. VIRTUAL ACCESS POINTS FOR VANETS

3.1 The protocol

A key obstacle to ubiquitous connectivity is that areas not covered by roadside APs cause intermittent access to mobile nodes. The proposed technique decreases the impact of such areas while being non-intrusive. Virtual Access Points (VAPs) are mobile nodes that act as a virtual AP broadcasting a series of objects or packets, based on their cache, when they move within the limits of a predefined region, outside any AP. In this way, pushed information from VAPs does not interfere with AP traffic. The role of a VAP is illustrated in Figure 1. It transparently re-broadcasts messages received from APs acting as an AP.




Figure 1 - A road coverage vision

The region where nodes act as a VAP can be either dynamically assigned or fixed and should be chosen depending on the areas covered from the nearest APs. This area defines the core of a VAP. The area lying in range of transmissions originating from the core VAP defines its coverage area. Dynamic VAP

positioning is more flexible, but long term assignment is easier to implement and is potentially closer to the optimal allocation.

Nodes are allowed to transmit only when they are in the core area. Any node passing through the core of a defined area, transmits only if there is no other node acting as VAP on that region. Figure 2 illustrates these concepts. It depicts an AP and a VAP. Two mobile nodes are currently connected to the AP, and receive information directly from it. Nodes 1 and 4 are in the core VAP area denoted with (a) and nodes 2, 3 lie in the coverage area of the VAP. Node 1 is currently acting as VAP, broadcasting its program. Although nodes 2 and 3 are both within range of node 1, they are “hidden” from each other. Therefore, if they start to transmit, they will not be able to detect each other’s transmission and messages sent will result to collisions at nodes 1 and 4. This is an example of what is called the “hidden terminal” problem. By adopting the notion of a core area for each VAP, with diameter roughly equal to the transmission radius, we avoid the “hidden terminal” problem, once each VAP candidate will, necessarily, be aware of every other candidate at the region. With this we also decrease unnecessary interference.




Figure 2 – A more detailed view

To efficiently utilize a VAP, transmissions from the nodes located inside its core area should be optimized. Amongst these nodes, only one each time acts as a VAP and composes the objects in its cache to a program that is subsequently broadcasted. By allowing a single node to become a VAP for as much time as possible, we ensure greater diversity among the objects transmitted and are able to favor the transmission of frequent and fresh objects, thus increasing the value of transmitted information.

3.2 Analysis

We assume that there is a finite set of objects, and that each time, each mobile node needs a subset of them. Such a subset can be the empty set or contain a number of objects, where each of them is associated with a deadline specified by the node. The deadline is the time after which the object is no longer needed and the query is dropped. When the node passes through roadside Access Points, it can query for objects and receive them with the answer, as well as overhear answers to queries of other nodes and objects that are proactively pushed by the AP. A cache is used to temporarily store objects for future queries. Whenever a node acts as VAP, it forms a program of cached objects to push to nearby vehicles.

The system works in a best-effort manner, and there is no guarantee that the VAP will help nodes to receive all the

messages they were meant to receive. However, the nodes acting as VAP select from their cache the objects most likely to be needed by other nodes.

In terms of exchanged data, VAPs monotonically increase the number of objects delivered to nodes. They operate in areas outside any APs and create no significant inference. The upper bound increase in the number of sent messages is given by:

$$nEM = nVAP * BS \quad (1)$$

where nEM is the number of exchanged messages, nVAP is the number of Virtual Access points, and BS is the Buffer Size. Unfortunately not all received messages are useful for every node and the duplicate or old objects are discarded. The number and locations of the VAPs will greatly affect the system’s performance. The data and latency of messages broadcasted on each a VAP depend on the distance from the originating APs.

The location of VAPs is also a sensitive parameter. A great concentration of VAPs near APs helps the dissemination of new information that the mobile nodes just received from the AP. Nodes moving towards an AP and passing through a VAP will receive messages that this AP has recently transmitted by nodes exiting the AP, and thus the “effective” range of this AP will be extended. However, in VAPs situated far from any AP, when the role of VAP is assumed by nodes that have not recently passed through any AP, the objects transmitted will be older, possibly less useful to receiving nodes. The further from APs a node is, the greater is the probability it will receive old and duplicated data, the points in-between APs being the most likely candidates for such areas.

We used the method proposed in [9] to develop and evaluate our technique. Our goal was to verify that the protocol is loop free. Even though it seems at first sight that our protocol is not loop free, we found that a loop scenario may occur when the relative velocities of the nodes are not equal. For example, considering the Figure 1, the simplest loop scenario occurs in the following case: node (A), acting as VAP in the point one, transmits the message M1 that is received by the node (B). Considering node (B) faster than node (A) and starting to act as a VAP at point two, it can transmit message M1, received by node (A). This characterizes a loop, and is one of the reasons why messages need to be equipped with unique IDs. Once the node (A) receives a duplicated message, identified by the ID, the node discards the message, indeed preventing the loop formation.

Another kind of message loop is present, and in fact is even desirable. Again, consider Figure 1, take the node (A) acting as a VAP in the lane 1, the message M1 sent can reach the node (C), going in the opposite direction in the lane 2. At some point in the future the node (C) start to act as a VAP and retransmits the message M1 that is received by the node (D) in the lane 1. If node (D) does not have the message, it is stored and will be retransmitted in the future incase node (D) becomes a VAP. However notice that this case is not a loop in the conventional sense, once the nodes involved are different. Other point to observe is that this kind of loop is even desirable once it helps spreading messages over the region. The buffer favors newer messages, so when older messages will be ignored and removed from the buffer.

The protocol's description allows for concurrent transmission. It is possible for two nodes to start acting as a VAP at the same time, when they are passing through the same region. This could become a problem since VAPs send messages using broadcast and the signal of the two VAPs will interfere with each other. For the simulation results we consider that the MAC layer's mechanisms handle this, either using CSMA-CA protocol, like the one described for the IEEE 802.11 networks, or a scheduler, as it is the case for IEEE 802.16 networks. In any case the worst thing that will happen is a waist of bandwidth due concurrency or even collision.

As stated before, the technique is a best effort kind. There is no guarantee the mobile nodes will receive all the messages needed to fill their buffers, or even that they will receive any message at all. It may happen that a node traverses the entire path from one AP to the other without receiving any message from other VAPs. This may happen in case the node is unfortunate enough to not be inside the VAP range of other nodes acting as VAP, or when the node itself is acting the VAP, and thus is not receiving messages from other VAPs. These situations are more likely to occur in sparse networks.

4. EXPERIMENTS

The scenario used on all the experiments to validate the technique consists of a road five Km long having four lanes, two in each direction. We have 33 different scenarios of five simulation minutes, with 200 vehicles on the road, roughly one vehicle at each 50 meters on average, at each direction. The vehicles max speed is 110Km/h and the min speed is 60 Km/h. The scenarios follow a realistic mobility pattern generated with the VanetMobiSim [10] tool. Each one of the generated scenarios have also 10 APs placed randomly along the five Km road. All results in the graphs show the confidence interval of 95%.

We used Sinalgo, a simulation framework for testing and validating network algorithms developed in Java by the Distributed Computing Group at ETH Zurich.

We were interested in the efficiency of our solution in two kinds of environments: by using a single access point and surrounding it with VAPs we wanted to test the technique "in vitro", to see how the system is affected by adding VAPs, while in the second case we cared more to see how VAPs would enhance an existing system with uncovered areas. Figure 3 illustrates these two separate cases. To determine the frequency by which each object was requested, we used an exponential mass function, an instance of which is depicted by figure 7.




Figure 3 – The two scenarios tested

Figure 4 depicts the results in the first case, where just one AP was present. The ratio between delivered queries and the total amount of queries is increasing, while that of dropped queries is

decreasing. As more VAPs are added, their relative contribution to the system drops. Figure 5 shows the same behavior for the case where we have more APs. In this case however, each VAP added is located between different APs and there is no apparent saturation. Figure 6 depicts the ratio between dropped and answered queries in the second case. We can observe that the ratio of answered queries, done, is consistently increase with the number of VBSSs. Queries that were never answered are considered dropped, while the ones whose response arrived after the deadline are considered neither done nor dropped.




Figure 4 – Efficiency of a single AP




Figure 5 – Efficiency for a road of 10 APs




Figure 6– Efficiency for a road of 10 APs




Figure 7– Frequency of Objects

5. Conclusions

This work introduces a simple yet powerful technique to increase network coverage for VENET networks. The technique is effective in increasing the packet reception rates with effectively no overhead on the AP regions. The next step for this technique is adding a pull based mechanism to permit to nodes to ask for information even when they are outside the real AP's range. VAPs are transparent to APs, and avoid causing interference by fixing the location where Virtual Access Points can transmit.

Extending VAPs in the case of streaming appears straightforward and we plan to measure the system's efficiency for streaming applications in the future. Furthermore, it would be interesting to consider which algorithms can be used for optimal VAP placement, and how the system's parameters should change when the vehicular traffic varies.

We also consider which algorithms will allow us to extend VAPs to enable a hybrid push-pull approach for VAPs and how this would affect the system's performance.

Finally, because the proposed VAP mechanism can be transparent to APs, it can be used as a subsystem of a larger system that would seek to offer connectivity in a world of almost ubiquitous wireless networks.

6. References

- [1] S. Das, A. Nandan, G. Pau, M.Y. Sanadidi and M. Gerla, SPAWN: Swarming Protocols for Vehicular Ad Hoc Wireless Networks, Proceedings of the First ACM International Workshop on Vehicular Ad Hoc Networks (VANET 2004), MOBICOM 2004, Berkeley, 2004
- [2] A. Nandan, S. Das, G. Pau, M.Y. Sanadidi and M. Gerla --- Cooperative Downloading in Vehicular Ad Hoc Wireless Networks, Proceedings of IEEE/IFIP International Conference on Wireless On demand Network Systems and Services, St. Moritz, Switzerland, Jan 2005
- [3] R. Shah, S. Roy, S. Jain, and W. Brunette. Data mules: Modeling a thre-tier architecture for sparse sensor networks. *Sensor Network Protocols and Applications*. IEEE, 2003.
- [4] Goodman, D.J., Borras, J., Mandayam, N.B., Yates, R.D., INFOSTATIONS: a new system model for data and messaging services, IEEE 47th Vehicular Technology Conference, Phoenix, AZ, USA, 1997
- [5] Zong Da Chen, H.T.Kung, and Dario Vah. Ad hoc relay wireless networks over moving vehicles on highways. In *MobiHoC*, 2001.
- [6] Qun Li and Daniela Rus. Sending messages to mobile users in disconnected ad-hoc wireless networks. In *Proceedings of the sixth annual international conference on Mobile computing and networking*, pages 44–55. ACM Press, 2000.
- [7] Briesemeister, L. and Hommel, G., Role-Based Multicast in Highly Mobile but Sparsely Connected Ad Hoc Networks, in *Proceedings of the First Annual Workshop on Mobile Ad Hoc Networking & Computing (MobiHOC)*, Boston, MA, USA, August 2000
- [8] Gavrilovich, C. D., *Broadband Communication on the Highways of Tomorrow*, in *IEEE Communications Magazine*, April 2001
- [9] D. Câmara, A. A. F. Loureiro, F. Filali, Methodology for Formal Verification of Routing Protocols for Ad Hoc Wireless Networks, IEEE GLOBECOM 2007, Washington, DC, November, 2007
- [10] J. Härri, M. Fiore, F. Fethi, and C. Bonnet, VanetMobiSim: generating realistic mobility patterns for VANETs, in *Proc. of the 3rd ACM International Workshop on Vehicular Ad Hoc Networks (VANET'06)*, September 29, 2006, Los Angeles, USA.
- [11] Swarup Acharya, Michael J. Franklin, and Stanley B. Zdonik. Balancing push and pull for data broadcast. In *ACM SIGMOD*, pages 183–194, 1997.
- [12] Dimitrios Katsaros and Yannis Manolopoulos. Web caching in broadcast mobile wireless environments. *IEEE Internet Computing*, 08(3):37–45, 2004.
- [13] Chi-Jiun Su and Leandros Tassiulas. Joint broadcast scheduling and user's cache management for efficient information delivery. *Wirel. Netw.*, 6(4):279–288, 2000.
- [14] Qiu Fang, Susan V. Vrbsky, Yu Dang, and Weigang Ni. A pull-based broadcast algorithm that considers timing constraints. In *ICPPW 04*, pages 46–53, Washington, DC, USA, 2004. IEEE Computer Society.
- [15] Liviu Iftode, Tamer Nadeem, Pravin Shankar. A comparative study of data dissemination models for vanets. In *Proceedings of the 3rd Annual International Conference on Mobile and Ubiquitous Systems: Networks and Services (MOBIQUITOUS 2006)*, July 2006.
- [16] Nikolaos Frangiadakis and Nick Roussopoulos. Caching in mobile environments: A new analysis and the mobile - cache system. In *IEEE International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC)*, Athens, Greece, September 2007.
- [17] M. Tariq, M. Ammar, E. Zegura. Message Ferry Route Design for Sparse Ad hoc Networks with Mobile Nodes. *ACM Mobihoc 2006*. May 2006 22-27, Florence, Italy.
- [18] P. Juang and H. Oki and Y. Wang and M. Martonosi and L. Peh and D. Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebraNet. *ASPLOS*, San Jose, CA, Oct 2002