VANET: On Mobility Scenarios and Urban Infrastructure. A Case Study

Gustavo Marfia*, Giovanni Pau*, Eugenio Giordano‡§, Enzo De Sena†, Mario Gerla*
*Computer Science Department - University of California, Los Angeles, CA 90095
e-mail: {gmarfia|gpau|gerla}@cs.ucla.edu
†Dipartimento di Ingegneria Elettronica e delle Telecomunicazioni, Universita‘ degli Studi di Napoli Federico II - Italy 80125
e-mail: enzodesena@gmail.com
‡Dipartimento di Ingegneria Elettronica Informatica e Sistemistica, Universita‘ di Bologna - Italy 40136
§Istituto Superiore Mario Boella - Italy
e-mail: egiordano@deis.unibo.it

***Abstract*— In [1] we show how vehicles can opportunistically exploit infrastructure through open Access Points (APs) to efficiently communicate with other vehicles. We also highlight the importance of the use of a correct mobility model, since the advantages that may derive from the use of an infrastructure may not be appreciated because of a the lack of accuracy.**

We continue our study based on realistic vehicular mobility traces of downtown Portland, Oregon, obtained from extremely detailed large scale traffic simulations performed at the Los Alamos National Laboratories (LANL). This mobility model is used to evaluate both flat and opportunistic infrastructure routing. We here build upon [1] and extend that work to: (a) assess the impact of a range of mobility models on network performance and; (b) discuss the performance trend we may expect during the day, as urban mobility patterns change.

We here compare results obtained with CORSIM [2] traces and Random Waypoint (RWP) [3] to the results obtained with realistic mobility traces.

I. Introduction

The setup of the UCLA Campus Vehicular Testbed (CVeT) [11], strongly motivates us in understanding the problems connected to the deployment of a Vehicular Ad-Hoc Network (VANET) with real traffic scenarios. The CVeT testbed will be composed of vehicles with both “periodic” (e.g. buses) and “random” (e.g. private cars) traffic patterns, thus stressing the network performance. An insight on real mobility patterns may already be drawn from a number of running testbeds (e.g. DieselNet [12]), but most of these experiments are built focusing on delay-tolerant application studies.

We here investigate the scenario where the great majority of cars are capable to connect to the network, the Dedicated Short Range Communications (DSRC) [10] initiative sets the path in that direction. We are interested in predicting the performance of those applications that require a real-time and a near real-time service. *Infotainment* applications, such as video streaming [15], gaming and peer-to-peer [14], are the expected first runners on a future vehicular grid.

The performance of a network, and, therefore of the applications that run on it, can heavily vary under different traffic conditions. In [1] we show that CORSIM traces, with a good level of detail, but with some lack of information in building the traffic model, can produce traces that are distant, in terms of network performance, from realistic traffic traces. Network performance problems that may be seen running a protocol on realistic mobility traces may not be appreciated using less accurate traces. In general, results wildly change from a mobility pattern to another.

A first vision of the role of the infrastructure in a vehicular grid may be found in [7]. In [1] we extend that work in new directions that aim to evaluate the feasibility of the vehicular grid, both with and without infrastructure, studying the impact of a realistic mobility model. We here produce a detailed feasibility study for the deployment of a public VANET in Portland, Oregon. We both evaluate the feasibility of a VANET seen as the wireless extension of the Internet through open APs and cars, and; compare the performance of routing protocols under both RWP mobility, CORSIM traces and realistic vehicular traces (LANL's traces) at different times during the day. Moreover, we aim to understand how CORSIM can be “tuned” to match realistic traces in terms of communication protocols assessment. Extensive work is available in both synthetic and trace driven simulations, but to the best of our knowledge, this is the first work that: (a) assesses, with realistic traces, the use of a real open AP infrastructure as a VANET infrastructure over a day long period (and, therefore, with very different traffic patterns); (b) attempts to analyze, through realistic traces, how a micro-simulator as CORSIM may be setup to produce sound traces for protocol evaluation. To reach these objectives, we use: (a) a realistic urban mobility model and; (b) a realistic infrastructure. We find (a) in the TRANSIMS mobility traces [6] and (b) in the open AP infrastructure of downtown Portland [9]. TRANSIMS realistic traces are first

Reference Author: Gustavo Marfia, Computer Science Department, University of California Los Angeles, CA 90095, e-mail:gmarfia@cs.ucla.edu
This work is partially supported by the Italian Ministry for Research via the ICTP/E-Grid Initiative and the Interlink Initiative, the National Science Foundation through grants Swarms and Whynet, and the UC-Micro Grant MICRO 04-05 private sponsor STMicroelectronics. Eugenio Giordano is supported by the Istituto Superiore Mario Boella, under the CVeT joint research lab initiative [11] funded by the Italian Ministry for Foreign Affairs and the Italian Ministry of Research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of our sponsors.

Fig. 1. *Street map of downtown Portland.*

found, in a networking paper, in [8], where the authors asses the performance of a large scale urban sensor network.

A number of mobility schemes, both trace driven (produced with a vehicular traffic simulator) and synthetic (model/equation based), have been proposed in the past few years. The availability of detailed street maps such as the TIGER database [4] and of commercial and affordable vehicular traffic micro-simulators such as CORSIM have driven the transition from simplistic synthetic models such as RWP and Constrained Random Waypoint (CRWP), to trace-driven, closer-to-reality models.

The paper is organized as follows. In Section II we describe the mobility models that are used in the paper. We then explain the simulation setting that will be used in Section III, where the results will be presented and commented. We finally conclude in Section IV.

II. Background

A. TRANSIMS Traces

We here build on the realistic traces drawn with TRANSIMS, a large scale vehicular traffic parallel simulator, based on activity flows.

An activity is the typical daily behavior of a household in a certain area. This information is collected through surveys and census data.

With this knowledge, it is possible to derive which are the typical movement patterns. In fact, from a large scale survey, it is possible to infer statistically sound schedules for the population set, and hence for the vehicles (nodes). In brief, the input to the simulator is the average behavior of a neighborhood household. Business sections, for example, are distinguished from residential areas, thus producing different traffic patterns. The TRANSIMS micro-simulator leverages on this information and builds a traffic model where the behavior of one cell is influenced by the behavior of neighboring cells (just as real traffic is). The simulator produces traffic traces which are tied to the node activities (for example, this car left home in a residential area at 6AM and got to the office across town, in the business section, at 6:45AM following a specified path).

B. CORSIM Traces

CORSIM is a vehicular traffic micro-simulator, which enables us to produce our own vehicular traces. The CORSIM simulator requires as an input: (a) a detailed map of the roads in the area, including speed limits (we find this information in the TIGER database); (b) the flow of cars per hour at each road that is cut by the map edges; (c) yield and stop signs; (d) traffic lights and their timing.

This micro-simulator is clearly not able to handle as many vehicles as TRANSIMS, since it runs on a single CPU. It also lacks activity information we don't have in running the simulations. The big advantage of using this tool is that we are able to evaluate, from a communication network standpoint, the effect of various levels of detail of the city map and of traffic flows. The CORSIM traces we use in [1] lack of traffic light timing information. Nevertheless, the average number of nodes and the average speed over all cars is the same. The only relevant difference, as may expected, is found in the average stop time. We observed this difference heavily influences the results in terms of delivery ratio, both with and without AP infrastructure.

III. Evaluation

A. Mobility Models Comparison: Simulation Setting

VANET simulations are run for 200 seconds on a 1 x 2 km rectangle on the map. The area has the highest AP density we found in Portland. It is highlighted in Fig. 1, located below the river and between the river and the highway.

Simulations are run in Qualnet [13]. Each simulation is set to end 10 seconds after the end of the last connection, so that no packets are still traveling at the simulation end. In the first 200 seconds timeframe, which starts at 7AM, we have an average of 270 vehicles at each time, an average speed of 12.6 meters per second (mps) (45kmph) and an average car stop time of 3.2 seconds. Starting at 8AM, there is an average of 371 vehicles at each time, an average speed of 12.5mps (45kmph) and an average stop time per car of 5.7 seconds.

We compare the delivery ratio as the fraction of sources increases. Sources and sinks are chosen at random and are initialized so that the fraction of active nodes remains constant on average. Each source sends a 20 seconds 4kbps CBR flow to a sink. 802.11b with auto-rate fallback is used at the MAC layer and Qualnet's two ray propagation model with shadowing is selected to simulate the wireless channel. The transmission power and the receiver sensitivity are set to reach a maximum transmission range of 250 meters. Finally, we plot the delivery ratio versus the average number of nodes that transmit, as the number of transmitting nodes increases. The graphs show the results obtained with up to 12% of nodes behaving as sources (i.e. we will have up to 12% of nodes sending and to 12% of nodes receiving in the area).

On such area we compare the performance of AODV [5] with TRANSIMS mobility with its performance with (a) CORSIM and; (b) RWP.

We will here show the results with two CORSIM runs. The first run is compared with TRANSIMS traces at 7AM, the second with TRANSIMS traces at 8AM. In both runs the average number of nodes and the average speed is close to the values at the same time of the day in TRANSIMS. The detailed values at 7AM are 277 average nodes, 10.4mps average speed

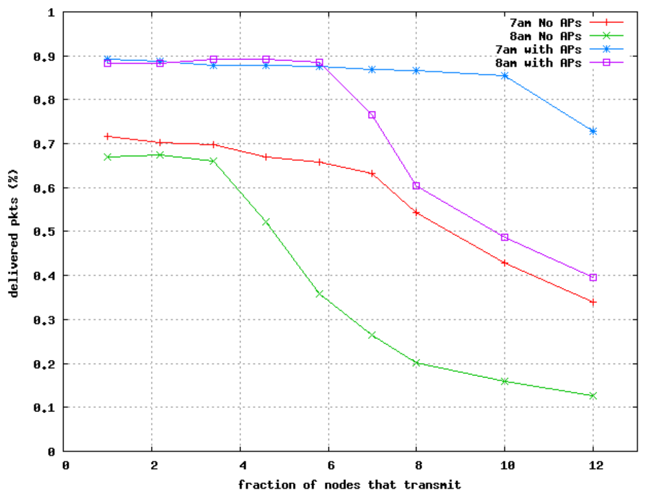

Fig. 2. *7AM and 8AM delivery ratio, with and without infrastructure, using the TRANSIMS mobility traces, as the number of sources increases.*

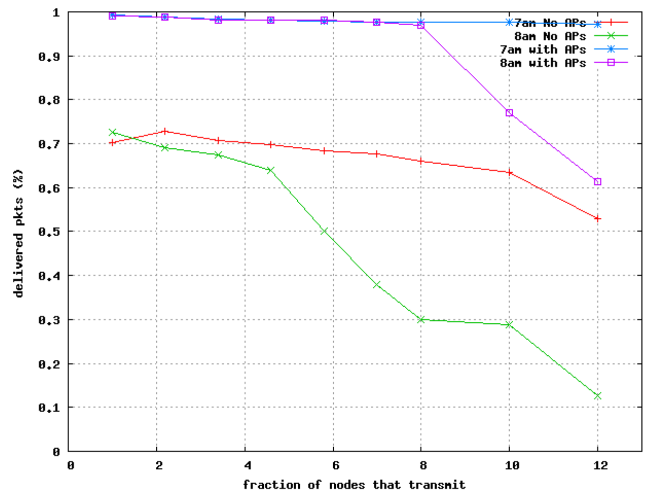

Fig. 3. *7AM and 8AM delivery ratio, with and without infrastructure, using the RWP model, as the number of sources increases.*

and an average stop time of 9 seconds. At 8AM the values are 377 average nodes, 10.4mps and 9.5 seconds stop time. In both runs we input the most relevant traffic lights in the area and their timing. The traffic light timing is derived from TRANSIMS. From one simulation to the other we change the uniform number of cars per hour flowing at each input street. At 7AM, at input streets on the map, we have an average flow of incoming and outgoing 55 cars per hour. At 8AM this value jumps to 85 cars per hour.

We finally compare the results with RWP mobility, where we tune RWP to match the average number of nodes, the average stop time and the average speed. For this reason we create two RWP models, one which we compare to the 7AM TRANSIMS run and one which we compare to the 8AM TRANSIMS run.

At 7AM and at 8AM we are in two different phases just before the rush hour. We can observe how results change with different car densities. We build on such results and observe the vehicular traffic trend in the area. In particular, we see how the number of nodes and the average number of neighbors influences the network's performance. As in [1], we view the effects of the opportunistic infrastructure on packet delivery ratio.

B. Mobility Models Comparison: Flat Network

Figs. 2, 3 and 4 show the results in terms of delivery ratio for AODV, at 7AM and 8AM, with the different mobility models. As already observed in [1], under TRANSIMS mobility, at 8AM performance initiates an abrupt breakdown beyond 4% of transmitting nodes. Overhead traffic congestion in this scenario becomes an issue, this reaches the 60% of the total traffic load in the network with the 12% of sources, thus collapsing the network. On the other hand, with the 7AM mobility traces, the breakdown point shifts to the 7% load point and beyond performance degrades more smoothly. The average speed of nodes in the two traces is very similar, clearly, the source of performance degradation does not reside in a higher mobility pattern. We then deduce that the reason of such different

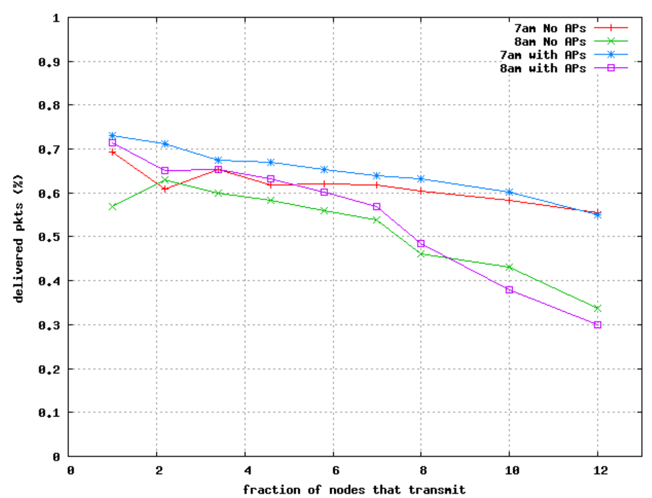

Fig. 4. *7AM and 8AM delivery ratio, with and without infrastructure, using the CORSIM mobility traces generated introducing traffic lights, as the number of sources increases.*

behavior mainly depends from the different density of cars in the different areas of the map.

The most surprising point in this result is it's similarity with the result we show in [1] for CORSIM traces with no traffic lights. We identified the performance breakdown with TRANSIMS mobility to reside in the high density of cars in certain areas of the map, due to vehicular traffic congestion. Introducing traffic lights in CORSIM simulations we expected to observe a similar breakdown pattern in delivery ratio for an increasing traffic load. This does not happen. We infer that we need to further refine CORSIM inputs to produce a mobility model that presents the same effects as those observed with TRANSIMS mobility.

RWP results are shown in Fig. 3, to complete this first comparison. Incredibly enough, RWP results are closer to TRANSIMS results, than CORSIM results. Nevertheless, RWP may not be considered representative for such type of study. With RWP, the breakdown point shifts to 6% and the delivery ratio is overestimated as the load increases. Moreover, due to the lack of topological information in RWP mobility, this similarity may not be generalized to other cases. In general it has been recognized that RWP poorly suites the need of foreseeing network performance in VANETs.

C. Mobility Models Comparison: Opportunistic Infrastructure

The 11 APs, found in [9], are inserted in the area and function as an infrastructure to the vehicular grid. We assume APs are connected in a star network configuration, with infinite bandwidth between them, so that changes in topology between the APs may be considered in-influential on the routing of packets between them. The assumption is consistent having that the open APs are connected to the fiber-to-the-home backbone, once packets are received by an AP they will flow over the fiber backbone.

We implement a two level hierarchical routing scheme, where the path between two endpoints (i.e. two cars) may or may not traverse APs. The objective of the routing scheme is to minimize the number of wireless hops traveled between two nodes. If, for example, the scheme should choose between two paths, where, say, the direct path (i.e. not including an AP) is made of 6 wireless hops and the indirect path, which involves AP traversal, includes 5 wireless hops, the routing scheme would clearly opt for the indirect path. We will then find two types of routes, the first type totally wireless, the second type wireless-wired-wireless. No wireless-wired-wireless-wired-wireless situation will be found, because of the infinite bandwidth assumption between two APs.

The heuristic rule we implement in the higher hierarchy level is to minimize the *wireless distance* traveled in a route. We here assume that a source knows the geographic coordinates of the destination and of the APs in the area and can therefore decide which is the best path (i.e. which may or not involve AP traversal) based on Euclidean distances.

At the lower hierarchy level we keep using AODV, which will attempt to find the shortest path by choosing the fastest return path. Based on the decision taken at the higher hierarchy level packet will flow, either directly from source to destination or from source node to AP1 and from AP2 to destination node.

Our previous work [1] shows that the opportunistic AP infrastructure can improve the performance of the vehicular grid. In that work we also show that, for the 8AM case, this improvement may have not been foreseen by using CORSIM traces produced with limited traffic and signaling information.

At 8AM, with TRANSIMS mobility, by simply exploiting open APs, the performance of the vehicular grid improves (from 70% to 90% of delivered packets) and the percentage of supported connections doubles (from 4% to 8% limit). At 7AM the performance improvement, in Figs. 2, confirms what observed at 8AM. By utilizing APs 10% of cars are able to transmit (i.e. 20% are involved in a communication) and delivery ratio ranges between 80% and 90%.

At 8AM the vehicular grid presents a higher density of cars (i.e. 100 more cars in the same area) and an almost twice average stop time (i.e. 5.6 vs. 3.2 seconds) than at 7AM, thus enabling cars to exploit the open AP infrastructure for longer periods.

Little changes are observed in using the open AP infrastructure in CORSIM, as may be seen in Fig. 4, for low loads. While we may have expected this to happen with no traffic lights, as we see it happens in [1], we would have expected a different behavior in this simulation, where traffic lights are introduced. With higher loads we see a 15% shift in performance, both at 7AM and at 8AM. Even if far from foreseeing the network breakdown point and the performance shift with APs, this result shows that it is beneficial to exploit stop times at traffic lights as the network load increases.

RWP again proves to be a better performance predictor than CORSIM with traffic lights, in the sense that foresees a performance improvement at both low and high loads. As for the flat network configuration, the performance improvement is more than what we can observe with TRANSIMS.

D. Discussion

In this section we build on the results from the previous sections in order to understand what is the expected behavior of the network on a longer timeframe. We have examined two 200 seconds timeframes up to this point, one at 7AM and one at 8AM, but we did not give any insight on the trend we may expect during the day.

First of all, we observe from Figs. 5 and 6, which represent the probability density functions for inter-contact time between two nodes, that the vehicular grid is typically well connected. Both graphs are saying that there is more the 90% the probability that a node will discover a new neighbor within one second at each time. We remind that in the examined area we have an average of 277 and 371 cars in the two time frames.

Fig. 7, derived from TRANSIMS mobility traces between 12:00PM and 6:00PM. Between 12:00PM and 2:45PM we have a total number of cars that resembles the amount we observed at 7AM and at 8AM. After 2:45PM the total number of cars in the area and the average number of neighbors per node steeply increases. We can then expect that the inter-contact time for a car traveling in the area after 2:45PM will be lower than the inter-contact time we can infer from Figs. 5 and 6. With this simple analysis it is clear that connectivity will never be an issue.

Figs. 7 and 8 show that the number of cars radically increase, the average speed in the area decreases and the fraction of static cars tends to one. As we may have expected simply driving around an average american city, heavy traffic jams build up. In a heavy traffic jam situation the network may be modeled as a static network. We therefore find that the most interesting timeframes to study, from a mobility model point of view, are those where we can still observe vehicular traffic mobility.

We may summarize the result of this section as follows: (a) a correct model of traffic flows is important, mobility models should go beyond matching the average values of traffic which prove to be poor predictors for the breakdown point of the network and the effect of open APs; (b) in long timeframes during the day (i.e. rush hours), the vehicular grid becomes static and so we may focus our attention, from a mobility model point of view, to shorter periods in the day.

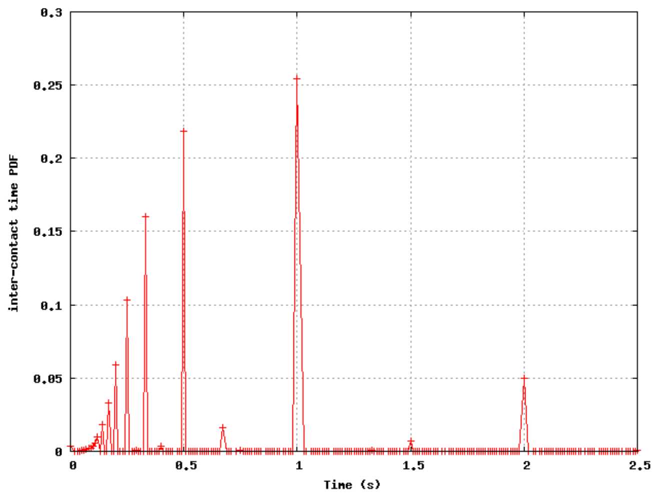

Fig. 5. *Inter-contact time, in TRANSIMS mobility traces, at 7AM.*

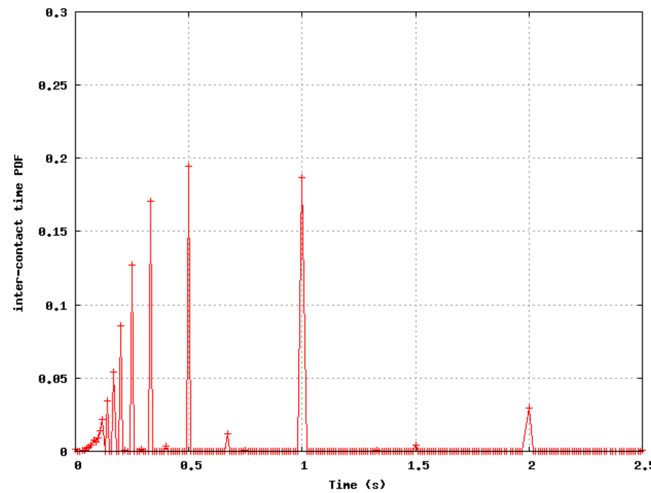

Fig. 6. *Inter-contact time, in TRANSIMS mobility traces, at 8AM.*

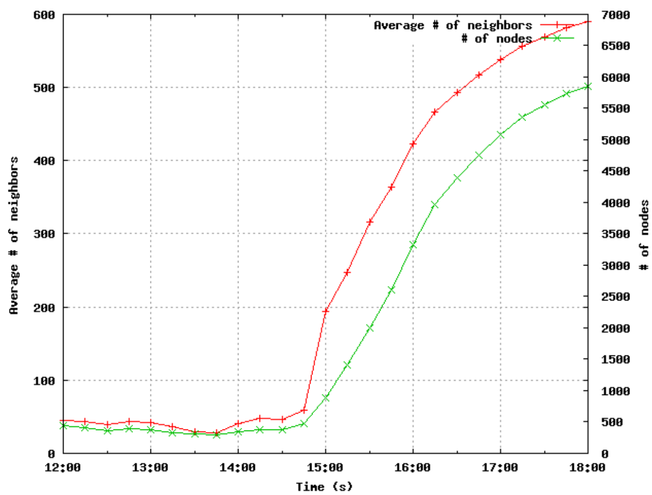

Fig. 7. *Total number of cars and average number of neighbors per car, between 12:00PM and 6:00PM, in the 1 x 2 km Portland area.*

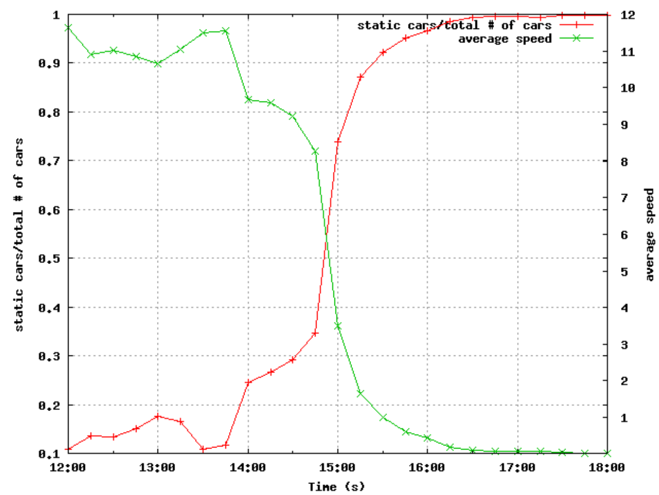

Fig. 8. *Average car speed and fraction of static cars, between 12:00PM and 6:00PM, in the 1 x 2 km Portland area.*

IV. Conclusion

This paper aims at giving an understanding of how a vehicular network could be setup starting from an existing infrastructure. We build on expand the work presented in [1] and show the effects of different mobility models and of different timeframes during the day on network performance. In particular, we consider a realistic case and examine a typical 7AM and 8AM traffic scenario in downtown Portland.

We confirm that open APs can and should be exploited and provide unexpected results. Performance improvements may be observed in both 7AM and 8AM results and we can foresee that as traffic congestion builds up, the importance of the role played by the AP infrastructure increases.

Synthetic models such as RWP are far from realism and bad performance predictors, for this reason part of our work goes in the direction of producing accurate CORSIM traces that may benefit the simulation of a realistic VANET. There are still some steps to take in this direction and are part of the future work in our research.

Acknowledgments

We wish to thank Dr. Stephan Eidenbenz and the Los Alamos National Laboratories for providing the realistic Portland traces this paper is based on.

References

[1] G. Marfia, G. Pau, E. De Sena, E. Giordano, M. Gerla, "Evaluating Vehicle Network Strategies for Downtown Portland: opportunistic infrastructure and the importance of realistic mobility models," in *The First International Workshop on Mobile Opportunistic Networking ACM/SIGMOBILE MobiOpp 2007, in conjunction with MobiSys 2007*, June 11, 2007, Puerto Rico, USA.

[2] [Available]. Online: http://mctrans.ce.ufl.edu/featured/TSIS/Version5/corsim.htm.

[3] J. Broch, D.A. Maltz, D.B. Johnson, Y.C. Hu, J. Jetcheva , "A performance comparison of multi-hop wireless ad hoc network routing protocols," in *Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, October 25-30, 1998, Dallas, Texas, USA.

[4] [Available]. Online: http://www.census.gov/geo/www/tiger/.

[5] C.E. Perkins, E.M. Royer, "Ad-Hoc On Demand Distance Vector Routing," in *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, pp. 90-100, February, 1999, New Orleans, Alabama, USA.

[6] [Available]. Online: http://transims.tsasa.lanl.gov/.
[7] B. Zhou, M. Gerla, F. Soldo, Y. Lee, G. Marfia, U. Lee, "Vehicular Grid Communications: The Role of the Internet Infrastructure," in *Wireless Internet Conference, Wicon 2006*, August 3, 2006, Boston, Massachusetts, USA.
[8] C.L. Barrett, S.J. Eidenbenz, L. Kroc, M. Marathe, J.P. Smith, "Parametric Probabilistic Routing in Sensor Networks," in *Mobile Networks and Applications*, Vol. 10, No. 4, pp. 529-544, 2005.
[9] [Available]. Online: http://www.auscillate.com/wireless/portland/.
[10] Online: http://grouper.ieee.org/groups/scc32/dsrc/index.html
[11] [Available]. Online: http://cvet.cs.ucla.edu/.
[12] [Available]. Online: http://prisms.cs.umass.edu/dome/.
[13] Online: http://www.scalable-networks.com/.
[14] A. Nandan, S. Das, G. Pau, M.Y. Sanadidi, M. Gerla, "Cooperative Downloading in Vehicular Wireless Ad Hoc Networks," in *Proceedings of Wireless On-Demand Networks and Services*, January, 2005, St. Moritz, Switzerland.
[15] M. Roccetti, M. Gerla, C. E. Palazzi, S. Ferretti, G. Pau, "First Responders' Crystal Ball: How to Scry the Emergency from a Remote Vehicle" , in *Proceeding the first IEEE International Workshop on Research Challenges in Next Generation Networks for First Responders and Critical Infrastructures (NetCri 07)*, April, 2007, New Orleans, Alabama, USA.