

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, Università degli Studi di Napoli Federico II - Italy 80125

e-mail: enzodesena@gmail.com

[‡]Dipartimento di Ingegneria Elettronica Informatica e Sistemistica, Università 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 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, gaming [15] 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 change 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. 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. In this paper we 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 found, in a networking

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.

paper, in [8], where the authors assess 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 which creates car movement patterns 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 the activity flow information we miss 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 as in TRANSIMS traces, over the same period. The only relevant difference, as may be 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




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.

values at 7AM are 277 average nodes, 10.4mps average speed 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 CORSIM runs we input the most relevant traffic lights in the area and their timing. The traffic light timing is derived from the TRANSIMS model. The TRANSIMS traffic inputs (to the area in question) are highly irregular. For simplicity, in CORSIM the input traffic is obtained by uniformly distributing the aggregate TRANSIMS input traffic over the boundary of the rectangular area shown in Fig. 1. As a result, at 7AM, at the streets entering and exiting the rectangle on the map, we have in CORSIM an average flow of 55 cars incoming and outgoing per hour. At 8AM this value jumps to 85 cars per hour. While these averages match the TRANSIMS model, the actual vehicle flow pattern is very different from one simulation model to the other, leading to strong discrepancies in performance, as we shall see later.

Next, we run a simulation experiment with RWP mobility, where we tune the RWP model to match the average number of nodes, the average stop time and the average speed measured in TRANSIMS. We create two RWP models, one for 7AM and another for 8AM traffic.

At 7AM and at 8AM we are in two different phases just before the rush hour. At 8AM the car density is much higher than at 7AM. 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 3.5% of transmitting nodes. Overhead traffic congestion in this




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.

scenario becomes an issue. The overhead 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 6% load point. Beyond this point, performance degrades more smoothly. The average vehicle speed in the two traces is very similar, clearly indicating that the source of performance degradation does not reside in a higher mobility pattern. We then deduce that the reason of such different behavior mainly depends from the different density of cars in the different areas of the map. A further investigation shows that most of the cars at 8AM are caught in a traffic jam in a small section of the map, near the freeway access. Consequently, most of the “randomly chosen” connections originate or terminate in this area. The “capacity” of this area is limited, thus an increase in offered load leads directly to a proportional increase in packet drop rate! At 7AM, there are less nodes and traffic is more uniformly distributed in the rectangle, so there is less

concentration of connections in the traffic jam area.

RWP results are shown in Fig. 3. RWP results are overly optimistic they overestimate network throughput. With RWP, the breakdown point shifts to 6% and the delivery ratio is overestimated as the load increases. RWP performs “random jumps” and thus does not accurately reproduce traffic, nevertheless RWP shows a similar trend to what we see with TRANSIMS model. We will further comment this in the last sub-section.

The CORSIM results are far from the TRANSIMS results, in spite of the fact that traffic lights and stop signs are carefully modeled to match TRANSIMS. The reason for the mismatch is due to the uniform input traffic assumption used in CORSIM. 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. Recall that we traced the performance breakdown of TRANSIMS to the high density of cars in certain areas of the map. The uniform input traffic assumption of CORSIM does not reproduce these hot spots; in particular it does not reflect the freeway on-ramp traffic jam observed in TRANSIMS. Consequently, the data traffic load is more uniformly distributed in the network. There is a gentle degradation in throughput as traffic load increases.

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 the specific topology connecting the APs has no impact on performance. The assumption is consistent with the fact that most public APs are connected with fiber links to the 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 path configuration needs to be explored, 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 those of the APs in its area. It 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 use AODV, which finds the shortest path either directly from source to destination or from source node to AP1 and from AP2 to destination node.

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). This improvement is due to the fact that the rerouting of traffic via the APs alleviates the traffic on the bottlenecks corresponding to the congested areas. 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%.

Improvements are observed in using the open AP infrastructure also in the CORSIM model, as may be seen in Fig. 4, for low loads and more significantly for high loads. In the latter case, we note a 15% shift in performance, both at 7AM and at 8AM. Still, the improvements are not as substantial as observed in TRANSIMS. In other words, CORSIM (with uniform traffic) is a conservative predictor. These results are as expected, as there is no traffic jam in the CORSIM model.

RWP proves to be an “unrealistically” optimistic performance predictor, showing strong performance improvement after AP rerouting at both low and high loads.

D. Discussion, Inter-contact Times and Afternoon Trend

From previous results, we see that CORSIM mobility is not producing the same network behavior as TRANSIMS mobility. Initially, by simply observing the average stop time of cars in CORSIM we would have expected a similar trend to what we observed with TRANSIMS mobility. A higher average wait time intuitively implies a higher density and an expected network breakdown due to congestion. We can explain this difference in the density structure of the traces. We observe that traffic, in the considered area, does not build up uniformly but rather back-pressures from the upper right corner of the map to all streets in the area. The upper right corner corresponds to a highway exit. At the peak of vehicular traffic congestion, all streets present a high car density. The CORSIM traces we produced lack into conveying the density area information, since, as already said, the flow rate has been set to be uniform at each incoming street on the map.

On the other hand the RWP model produces a mobility pattern where the greatest density of cars, with a higher probability, may be found at the center of the map [16]. The center of the map is only a few blocks away from where the highest traffic density is in TRANSIMS traces, at the highway entrance at the upper right corner of the map. We should also notice that the highest density of APs is in the left area. The high density area in RWP is closer, then, to APs, than the high density area in TRANSIMS. This explains why performance improves so much, with APs, with RWP mobility.

We now 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.




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.

Both graphs indicate there is more than the 90% chance that a node will discover a new neighbor within one second at each time. Please remind that in the examined area we have an average of 277 and 371 cars in the two time frames.

Fig. 7 is derived from TRANSIMS mobility traces between 12:00PM and 6:00PM. Between 12:00PM and 2:45PM we have a total number of cars comparable to 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, as 5PM rush hour approaches, 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 of the network breakdown point and of the effects 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 during the day.

IV. CONCLUSION

This paper expands on the work presented in [1] and shows the performance impact of different mobility models at different times during the day. Specifically, we examine a typical 7AM and 8AM traffic scenario in downtown Portland and then, based on those results and more data, extend the discussion to the afternoon.

The results confirm that open APs can be effectively exploited to dramatically improve performance, no matter which mobility model is used.

Besides the above “positive” result, which is anyway rather intuitive, another contribution of this paper is that of alerting us of the importance of accurate vehicle density modeling on performance. In our CORSIM model we took great care in matching traffic lights and stop signs to those found in the TRANSIMS reference model. However, for practical reasons we used a uniform vehicle traffic input model to our urban map. As a result, the distribution of vehicles in the urban map was not very accurate. Performance was grossly overestimated. RWP behaved a little better, but such result cannot be extended to a general VANET scenario, as we explained. In any case, RWP performance predictions were overly optimistic.

Future work will be directed to the development of synthetic models that generate more realistic traffic inputs. These models must faithfully reproduce the spatial correlations, vehicular densities and traffic jams experienced in real world situations. Once these “accurate” models are developed, they will be used to evaluate a broader family of routing protocols beyond AODV.

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.auseillate.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] C.E. Palazzi, “Fast Online Gaming over Wireless Networks”, PhD Thesis, UCLA, CA, USA, 2007.
- [16] G. Resta, P. Santi, “An analysis of the node spatial distribution of the random waypoint model for Ad Hoc networks,” in *Proceedings of ACM Workshop on Principles of Mobile Computing (POMC)*, pp. 44-50, October 2002, Toulouse, France.