Evaluating Vehicle Network Strategies for Downtown Portland: Opportunistic Infrastructure and the Importance of Realistic Mobility Models

Gustavo Marfia
UCLA
gmarfia@cs.ucla.edu

Giovanni Pau
UCLA
gpau@cs.ucla.edu

Enzo De Sena
Universita' degli Studi di Napoli Federico II
enzodesena@gmail.com

Eugenio Giordano
Istituto Superiore Mario Boella
Universita' di Bologna
egiordano@deis.unibo.it

Mario Gerla
UCLA
gerla@cs.ucla.edu

ABSTRACT

In an urban environment, vehicles can opportunistically exploit infrastructure through open Access Points (APs) to efficiently communicate with other vehicles. This is to avoid long wireless ad hoc paths, and to alleviate congestion in the wireless grid. Analytic and simulation models are used to optimize the communications and networking strategies. For realistic results, one important challenge is the accurate representation of traffic mobility patterns.

In this paper we introduce realistic vehicular mobility traces of downtown Portland, Oregon, obtained from extremely detailed large scale traffic simulations performed at the Los Alamos National Laboratories (LANL). To the best of our knowledge, these are among the most accurate synthetic motion traces available for study, with the exception of actual car trace measurements. The new mobility model is used to evaluate AODV [1] in flat and opportunistic infrastructure routing. To assess the importance of a realistic mobility model for this evaluation, we compare these results with those obtained with CORSIM [2] traces.

The paper makes the following contributions: (a) introduction of efficient, opportunistic strategies for extending the AP infrastructure to use vehicle to vehicle paths, and (b) assessment of different mobility models - CORSIM traces and LANL's realistic vehicular traces - in the modeling of different routing strategies.

Categories and Subject Descriptors: C.2.1 [Computer Communication Networks]: Network Architecture and Design

General Terms: Algorithms, Performance

Keywords: AODV, Ad Hoc Network, MANET, Opportunistic Networking, Portland, VANET, Vehicular

Figure 1: ***Street map of downtown Portland.***

1. INTRODUCTION

The problem of defining the best possible simulation scenario for a Vehicular Ad-Hoc Network (VANET) has been studied along the years. The prohibitive costs and the lack of repeatability of a real world deployment induce researchers to test protocols and traffic patterns in simulation. Performance results can widely vary from a static environment to a mobile one, and from a mobility pattern to another, for such reason the definition of the *right* realistic mobility model is a hot research area. A number of mobility schemes, both trace driven and synthetic, have been proposed in the past few years. More recently, the availability of tools such as TIGER [3] and CORSIM have driven the transition from simplistic synthetic models such as Random Waypoint (RWP) [4] [15] to, trace-driven, closer-to-reality models (at least in terms of spatial constraints on nodes).

The mobility traces we use are drawn from large scale simulations performed at LANL with TRANSIMS [5], a large scale distributed vehicular traffic simulator, based on activity flows. Differently from random based models, where the direction and speed of a node are drawn at random step by step, with or without memory and neighborhood awareness, the motion of a node through the urban grid is inferred from a survey. 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). By mixing this information and the cellular automata model of the TRANSIMS micro-simulator, 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 at 6AM and got to the office across town at 6:45AM following a specified path). We rely on such realistic traces for the evaluation of a VANET in downtown Portland and as a benchmark for simpler models.

The availability of realistic traces, such as the Portland traces, does not lower the interest in accurate synthetic models for vehicular traffic flows. The traces produced by LANL require extensive processing (grid infrastructure), statistically sound information (activity surveys) and require time and effort to generate. We here compare CORSIM traces, which are computationally feasible to produce, against the realistic traces by observing the behavior of the communication network. We use the AODV routing protocol, both with and without infrastructure. Previous work [6] shows that the performance of such routing protocol heavily degrades from

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

synthetic mobility models to realistic traces. We here observe that such difference heavily depends from the communication and vehicular traffic patterns.

We envision a VANET as a superposition of cars and open infrastructure [8] and underline the role of APs in its deployment. The strong development of city wireless (e.g. Google WiFi at S. Francisco) and of consumer electronics devices in cars brings such vision close in the future. Under such assumption, we think of a VANET rooted at APs that stretches through car multi-hop. For such reason, we investigated which are the available, up to date, open APs in Portland [10], found their coordinates, and considered their suitability as an infrastructure.

In this paper we put mobility models to test in a feasibility study for the deployment of a public VANET in Portland, Oregon. We will 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 CORSIM traces and a realistic vehicular trace based on activity location models.

Extensive work has been published on ad hoc net motion modeling. A survey may be found in [11]. TIGER maps are introduced to spatially restrain traffic flows to urban streets in [12]. A model that considers the state of neighboring nodes in the motion of a car is shown in [13]. The case of swarm and group mobility models, which may be best suited for military scenarios, is considered in [14]. Work that presents TRANSIMS traces, albeit for a different purpose than ours (i.e. sensor network deployment) may be found in [9], while the first large scale simulations for the Zurich case are performed in [6]. A first attempt to consider the impact of propagation loss in an urban scenario may be found in [7].

The paper is organized as follows. In Section 2 we describe the mobility models that are used in the paper and how these are compared. We then explain the simulation setting that will be used in Section 3, where the results will be presented and commented. We finally conclude in Section 4.

2. BACKGROUND

2.1 TRANSIMS Traces

The Downtown Portland mobility traces cover a 900 sec (15 minutes) time frame starting at 8AM. The traces provide the position of each vehicle (cars, buses, etc.) with a one second granularity. The considered area is 3 x 7 km wide, an average of 3,000 cars are seen at each time instant and total of 16,529 cars are seen through all the period. 2,891,257 are the total events (positions) recorded in the 15 minute trace.

Figure 1 shows that both urban and highway traffic is present in the area. We should also notice that, due to the geographical configuration of the city in the considered area, the city is split in two by the river, not only in geographical terms, but also in communication terms.

The car speed ranges between 0 and 37.5 meters per second (mps), or 135kmph, in 7.5mps (27kmph) steps. We appreciate the trace's realism by seeing that the great majority of events have a speed of 54kmph (urban speed), while highway events, plus perhaps some exceedingly fast city driver, is represented by the 81, 108 and 135kmph speeds. Only a small portion of cars are standing still (the 20% of events at 0kmph include cars waiting at traffic lights and stop signs), as we may expect at the beginning of the morning rush hour.

The number of nodes slowly increases between 8am and 8:15am, as we may expect in the morning in a typical American town, bringing the car density from approximately 143 cars per square km to 167 cars per square km. Accordingly, we also have an increase in the average number of neighbors per node, from 57 to 70 neighbors. Two nodes are defined to be *neighbors* if they are within a 250 meter range from each other.

As for Portland infrastructure support, we found 31 open APs in the area [10]. The position of the APs are represented by an X in Figure 1. The APs we found are typically located in hotels and cafeterias and serve the purpose of providing Internet connectivity to customers. This is clearly approximate information, APs can appear and disappear at a fast pace and the information used to locate the APs may be out of date at the time of this writing. We view this data, however, as representative and adequate to be used as the basis for our feasibility study. We will update the result as fresher inputs are obtained.

2.2 CORSIM Traces

As TRANSIMS, CORSIM enables to reach a high level of precision in vehicular traffic simulation. Information such as traffic lights timing and phase, street lanes, speed limits and car flows may be input. Differently from TRANSIMS, CORSIM runs on a single processor and lacks of the parallelization which would lead to large scale simulations.

We here extract maps and street information from the TIGER database and use it as an input for CORSIM. We tune CORSIM input parameters (i.e. car flow at a boundary street) to match the average number of cars observed in the area in TRANSIMS traces. Street speed limits are found in TIGER. Car speeds are automatically adjusted by the simulator by implementing normal traffic behavior. Because of the lack of traffic light timing information and stop sign location information, we did not provide the CORSIM vehicular simulator with such type of input parameters. We then experience a smaller average stop time in this set of traces, as cars find no constraints to stop them, other than avoiding to provoke a car accident.

With little information (i.e. street speed limits and the average number of cars in an area in a certain time interval), we produce mobility traces for Portland, Oregon.

3. EVALUATION

3.1 AP infrastructure

The use of a vehicular grid together with an infrastructure has already been discussed in [8]. We rely on a similar idea, analyzing a realistic case and its feasibility. We will here analyze what are the benefits of using the opportunistic infrastructure provided by the open APs, by evaluating the average hop distance of mobile nodes (i.e. cars) during the 900 sec time period.

Car and AP transmission range clearly determines the average number of hops between a car and an AP. AP positions also may influence the number of hops. To test this dependence we first use a highly abstract model that uses Euclidean distance (divided by range) to approximate hop distance. This is obviously a very crude estimate that provides a starting point of reference. Figures 2 and 3 show us how the average number of hops from a car to an AP varies depending on AP transmission range, car transmission range and position. In Figure 2 we place the real APs on the map and vary AP and car transmission range. We observe that with realistic choices of car and AP transmission ranges, 100mt for a car, 50mt for an AP (recall that APs are placed in buildings and their signal strength is attenuated by walls), the average car to AP hop count is 16. We then take the APs out of the buildings and place them randomly at street intersections. With a sample size = 50 trials and choosing a more optimistic transmission range for APs (i.e. 100mt) the random intersection placements reduced the average hop count

Figure 2: ***Average number of hops, from a node to one of the 31 APs on the map, as the car transmission range increases. We have two curves, one where AP transmission ranges are set to 50mt, the second where AP transmission ranges are set to 100mt.***

Figure 3: ***Each point in the graph is produced averaging over 50 trials. In each trial APs are placed at random on the map, following a uniform distribution. AP transmission range is set to 100mt. We plot average car-to-AP hop count vs. AP count on the map.***

from 14 to 7. With no doubt a more careful, traffic dependent AP location strategy would further improve this value.

In the previous figures we derived hop count rather simplistically. The value of 14 average hops from a car to an AP is clearly an optimistic lower bound. In order to have a more realistic assessment, we use the popular routing protocol AODV to compute path lengths. The estimate is still not very realistic as we use a static network with ideal channel (disk propagation model). The result is shown in Figure 4. As expected, more realistic AODV routing requires almost twice as many hops to destination in case, ideally, 7 hops separate two nodes. This can be explained by reminding that AODV will not necessarily pick the shortest path, but rather the fastest return path. Clearly, from these results, we should infer that the current AP Portland coverage is unsuited to support a vehicular grid. The vehicular grid can exploit opportunistically the infrastructure in a very limited manner for real-time and near real-time communication traffic. We will further evaluate this statement in the last subsection.

Figure 4: ***Comparison between the ideal number of hops between two nodes (distance divided by transmission range) and the number of hops computed between the same nodes by AODV.***

3.2 Mobility Models Comparison: Simulation Setting

Next, we evaluate the performance of AODV using the TRANSIMS traces, versus the performance using CORSIM traces. We also view the effects of the opportunistic infrastructure on packet delivery ratio. Because of the number of nodes involved in the simulation and the scaling problems of Qualnet, we selected a 1 x 2 km rectangle on the map and simulated the VANET in the selected area for 200 seconds. The chosen rectangle is the most dense area in terms of APs. The area is highlighted in Figure 1 and may be visually located as the section below the river and between the river and the highway.

The simulation time is set to end 10 seconds after the end of the last connection, so that no packets are traveling when the simulation stops. In the 1 x 2 km area, in the 200 seconds timeframe, we have an average of 371 vehicles at any one time, for an average speed of 12.5mps (45kmph) and an average stop time per each car of 5.7 seconds. We are in this way able to enter consistent parameters in the CORSIM simulator.

The average number and the average speed of nodes are the same in the CORSIM traces. The only difference is in the average stop time that is 0.2 seconds, lower than in TRANSIMS, since we didn't input traffic light timing to the CORSIM simulator.

We compare delivery rate and packet delay in uniform traffic, as the percentage of traffic sources increases. Nodes are chosen at random in the simulation area. The fraction of active nodes is constant on average. An active node is set to send a 4Kbps CBR flow to a random peer. Each connection lasts 20 seconds on average. The MAC layer we use is 802.11b with auto-rate fallback. We choose to use the two ray propagation model with shadowing, available in Qualnet. The transmission power and the receiver sensitivity are set to reach a transmission range of 250 meter. Finally, we increase the average number of nodes that transmit up to 12% of the nodes on the map (i.e. we will have up to 12% of nodes sending and to 12% of nodes receiving in the area).

3.3 Mobility Models Comparison: Flat Network

We here discuss the case where no AP infrastructure is opportunistically exploited by the vehicular grid. We will then compare and verify whether there is any significant advantage in using the available APs.

Figures 5 and 6 show the results in terms of delivery ratio for AODV, using the two mobility models. What here captures our attention is that with the TRANSIM traces performance experiences an abrupt breakdown around 4% of transmitting nodes. On the other hand, with the CORSIM mobility traces, performance degrades more smoothly. The overall average speed of nodes in the two traces is the same (by construction), however, the fact that our CORSIM traces do not include traffic lights and stop and go effects (while TRANSIMS traces do) leads to two important consequences. First, in TRANSIMS, stop signs and traffic lights, combined with morning traffic, highly increases the density of cars in specific areas of the map (e.g. busy intersections etc.) as opposed to others. Because of the overall high density of nodes in these areas we can then observed that overhead traffic congestion becomes an issue above a certain load with the TRANSIMS mobility pattern, causing degradation. Second, the stop and go behavior causes the distances between cars on the AODV path to fluctuate wildly, leading to path breakage. Each path break must be recovered with a flood search of a new AODV path. Both these effects, combined, cause a tremendous load on the system, which progresses linearly with the fraction of active nodes, causing severe congestion and eventually collapse.

In CORSIM, the path breakage is much less frequent because car maintain a steady distance from each other. In fact, because of our simulation settings, nodes are spread more evenly in the CORSIM traces. This explains the more gradual decrease of performance as the number of connections increases.

This result clearly shows the importance of accurate mobility models for VANETs.

3.4 Mobility Models Comparison: Opportunistic Infrastructure

We now insert the 11 APs that are in the rectangle. We connect the APs in a star network, where the bandwidth between APs is set to 100Mbps. We are assuming that the open APs are connected to the fiber-to-the-home backbone and that once packets are received by an AP they will flow over the backbone.

We here implement a two level hierarchical routing scheme, where the path between two endpoints (i.e. two cars) may traverse or not 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

Figure 5: ***Delivery ratio, with and without infrastructure, using the TRANSIMS mobility traces, as the number of sources increases.***

Figure 6: ***Delivery ratio, with and without infrastructure, using the CORSIM mobility traces, as the number of sources increases.***

hops, the routing scheme would clearly opt for the indirect path.

The heuristic rule we implement in the higher hierarchy level is to minimize the traveled *wireless distance*. We assume that a source has a knowledge of the positions of the destination and of the APs and is able to choose 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, based on the decision taken at the higher hierarchy level, either directly from source to destination or from source node to AP1 and from AP2 to destination node.

The question, now, is to understand how the vehicular grid is helped by the opportunistic infrastructure. Previous results show that flat routing poorly performs as the number of connections increases. With such low delivery ratio, no application may survive (unless some revolutionary coding scheme is invented in the meanwhile).

Figures 5 and 6 show the results in terms of delivery ratio. We observe an improvement in both cases, but the most important point is to observe how delivery ratio shifts with the TRANSIMS traces in Figure 5. Delivery ratio keeps a steady value around 90% up to 8% of sending nodes. In the interval 5-8% performance steps from as low as 5% up to 90% of packets delivered to destination. This is, in practice, an important result: with no effort, by simply using the infrastructure, the performance of the vehicular grid can be improved (from 70% to 90% of delivered packets) and the percentage of supported connections doubled (from 4% to 8% limit). We observe a slight improvement with CORSIM traces in Figure 6.

The reason for the shift of the collapse point for TRANSIMS from 4% to 8% load can be explained by the fact that the paths to APs are more stable than the inter-car paths, thus requiring a lower path reconstruction rate. This can be explained by the fact that paths through access points involved fewer wireless hops. Moreover, APs tend to be positioned in areas with heavy (and slow) traffic, high density and lower car to car distance fluctuations. These two effects contribute to the performance improvement.

Clearly, as we just pointed out, the performance is affected by the degree of overlap between the areas served by access points and the areas with greatest car density. The position of an AP in an area where cars stop is beneficial for the network. As cars stop, they have a greater chance of exploiting the infrastructure and reach a peer with fewer (and more stable) wireless hops. We can say that car traffic near open APs improves the overall delivery ratio of the network or, more precisely, APs placed near highly dense areas improve overall performance.

With the CORSIM traces, as previously discussed, we do not observe relevant differences in performance. Cars rarely stop and hence do not opportunistically exploit the infrastructure. This is the main explanation of the different pattern of performance improvement between the TRANSIMS and the CORSIM case.

4. CONCLUSION

This paper helps our understanding of how the infrastructure can help vehicular communications. We considered a realistic case and examined a typical 8AM traffic scenario in downtown Portland exploring both options with and without infrastructure. Two main results are reported: (a) the motion model has an enormous impact on performance - the CORSIM simulation (deprived of stop and go dynamic behavior cannot predict the sudden throughput collapse), and; (b) the presence of APs and infrastructure make a big difference in performance but again this difference can be appreciated ONLY with the accurate motion model.

A future study will evaluate the performance of CORSIM when the model includes also traffic lights and middle of the block acceleration and deceleration effects. We will compare the revised CORSIM with TRANSIMS. A second important future direction will be to compare the results we here show with results that consider other time frames during the day. For example, 8AM traffic is much different from 1PM traffic, and it would be interesting if the observations we here make may be generalized in time. A third direction should consider new mobility models, including synthetic mobility models (e.g. RWP , CRWP and Real Track models). A final point is to include new routing schemes and examine how they behave in the aforementioned scenarios.

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. 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 [16] 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.

5. REFERENCES

[1] 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.

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

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

[4] 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.

[5] [Available]. Online: http://transims.tsasa.lanl.gov/.

[6] V. Naumov, R. Baumann, T. Gross, "An Evaluation of Inter-Vehicle Ad Hoc Networks Based on Realistic Vehicular Traces," in *Proceedings of the the Seventh ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, May 22-25, 2006, Florence, Italy.

[7] E. Huang, W. Hu, J. Crowcroft, I. Wassell, "Towards Commercial Mobile Ad Hoc Network Applications: A Radio Dispatch System," in *Proceedings of the the Sixth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, May 25-27, 2005, Urbana-Champaign, Illinois, USA.

[8] 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.

[9] 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.

[10] [Available]. Online: http://www.auscillate.com/wireless/portland/

[11] T. Camp, J. Boleng, V. Davies, "A Survey of Mobility Models for Ad Hoc Network Research," in *Wireless Comm. & Mobile Computing* $WCMC$:

Special issue on Mobile Ad Hoc Networking Research, Trends and Applications, Vol. 2, No. 5, pp. 483-502, 2002.

[12] A.K. Saha, D. Johnson, "Modeling Mobility for Vehicular Ad Hoc Networks," extended abstract, in *VANET '04*, October 1, 2004, Philadelphia, Pennsylvania, USA.

[13] D.R. Choffnes, F.E. Bustamante, "An Integrated Mobility and Traffic Model for Vehicular Wireless Networks," in *VANET '05*, September 2, 2005, Cologne, Germany.

[14] B. Zhou, K. Xu, M. Gerla, "Group and swarm mobility models for ad hoc network scenarios using virtual tracks," in *Milcom '04*.

[15] D.B. Johnson, D. Maltz, "Dynamic source routing in ad hoc wireless networks," in T. Imelinsky and H. Korth, editors, *Mobile Computing*, pp. 153-181. Kluwer Academic Publishers, 1996.

[16] [Available]. Online: http://cvet.cs.ucla.edu/.