Feasibility Study of Using Mobile Gateways for Providing Internet Connectivity in Public Transportation Vehicles

Gunadi Setiwan, Samuel Iskander,
Salil S. Kanhere, Quan Jun Chen
School of Computer Science and Engineering
The University of New South Wales
Sydney, Australia
{gste047, samueli, salilk, quanc}@cse.unsw.edu.au

Kun-Chan Lan
National ICT Australia
Sydney, Australia
Kun-Chan.Lan@nicta.com.au

ABSTRACT

The extension of Internet services to public transport passengers is slowly becoming inevitable. Several architectures for providing Internet access to moving vehicles have been evaluated in the past. However, most of these studies have focused on using static gateways. In this paper, we study the feasibility of an architecture that involves deploying mobile gateways on a selected subset of public transport vehicles for providing Internet connectivity to the entire fleet. The vehicles organize as dynamic clusters and connect to the Internet by communicating with the gateways via multi-hop paths. We evaluate the underlying connectivity characteristics and the coverage achieved by employing an optimal gateway placement strategy. In our analysis, we use realistic movement patterns of public transport buses in a metropolitan city. We also propose a prediction based enhancement, which takes advantage of the known mobility patterns of the buses to improve the performance of the multi-hop routing protocols employed within each cluster.

Categories and Subject Descriptors: C.2.1[Network Architecture and Design]: Wireless communication; C.2.2[Network Protocol]: Routing protocols

General Terms: Algorithms, Measurements, Performance

Keywords: Mobile Gateways, Vehicular Ad-hoc Networks, Realistic Mobility Patterns, On-Board Communication

1. INTRODUCTION

Recently a new paradigm of *Networks in Motion* is quickly attracting interest from the research community and is also being viewed as a viable commercial solution [1], [2], [3]. A typical on-board network, as illustrated in Fig. 1, consists of an on-board LAN (wired and/or wireless), which is connected to the Internet through a mobile gateway. On-board users can simply plug in their devices to the on-board LAN and access the Internet. The gateway can utilize a diverse array of wireless access technologies (e.g: GPRS, UMTS, 802.11) through multiple service providers. However, deploying a gateway on each bus is quite expensive considering a large number of buses. In this work, we propose an architecture in which a small number of mobile gateways are deployed on a few selected buses. We assume that each bus is equipped with a wireless device which allows it to communicate with buses within its radio range and form a self-organized ad-hoc network cluster. While travelling along their regular routes, the buses dynamically join and leave clusters. The rest of the buses connect to the Internet by communicating with the gateways via multi-hop paths.

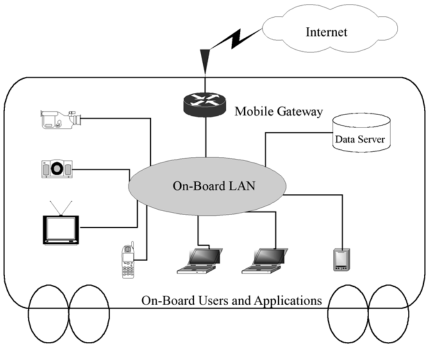

Figure 1: A typical on-board network

In this paper, our goal is to evaluate the feasibility of such an architecture, for which we use real-world mobility traces collected from the bus system of a metropolitan city. Moreover, we want to address the following questions:

- Given that the size of cluster is strongly affected by the radio range of each bus, which existing radio technology might be more feasible for such an architecture (Section 3)?
- How many mobile gateways need to be deployed for providing an acceptable level of connectivity to all buses (Section 5)?
- Which existing ad-hoc routing protocol is more suitable to be used in the ad-hoc clusters formed amongst the buses (Section 6)?

IWCMC'06, July 3–6, 2006, Vancouver, British Columbia, Canada.

Given that the mobility patterns of the buses are known apriori, we propose a mobility-related prediction-based strategy which allows the buses to predict the closest gateway within their multi-hop cluster and study its improvement on the routing protocols (Section 6).

The rest of this paper is organized as follows. Section 3 describes the traces and presents characteristics of its mobility patterns. Section 4 presents the metrics that we use to analyze the underlying connectivity achieved by our architecture. The evaluations and associated discussions are presented in Section 5. Finally, Section 6 presents our mobility prediction based enhancement and evaluates its effect on several routing protocols.

2. RELATED WORK

Namboodiri et al., [4] studied the feasibility of placing mobile gateways on selected vehicles to provide connectivity to the other vehicles in their vicinity. Their proposed architecture is similar to ours but there are several differences. Firstly, their simulations were conducted for a highway scenario wherein the nodes and mobile gateways were uniformly distributed along a straight long highway. Moreover, they assume a first-order Markov model to characterize the motion of the vehicles. In our work, we are using realistic movement traces of buses within a metropolitan city.

Huang et al. [5] proposed an application scenario for mobile ad hoc networks in the form of a radio dispatch system for taxis and investigate its financial and technical feasibility. In their evaluations they modeled the city as a grid [6] of size 5 km x 5 km, with 300 taxis distributed within this area. The coverage and outages experienced and concluded that their system does perform satisfactorily under most operating conditions. Their focus is however on studying the performance of an application on a purely ad hoc network, which is different from the architecture that we aim to evaluate in this study.

Ad Hoc City [7] present a multi-tier architecture for providing Internet connectivity to mobile users. In their architecture, several fixed Internet connected base stations are deployed throughout the city. A multi-hop network composed of wireless devices mounted on mobile vehicles such as cars and buses forms the backbone, and the vehicles connecting to the static gateways as they move through the city. Individual users can utilize the mobile backbone to access the Internet, with the backbone relaying the user packets to the static gateways, either directly or using multi-hop routing. They analyzed the performance of their architecture using real-world bus movement traces. In this paper we have used the same traces for evaluating our proposed architecture. However, unlike their scheme we do not rely on static gateways since there could be several outages in the coverage when the relaying nodes are disconnected from the static nodes. Furthermore, the cost associated with leased fixed base stations and the associated maintenance can be very high. Similar architectures have been discussed in [8], [9], wherein vehicles connect to the Internet as they drive-by static gateways, which are periodically deployed along the highways. However, this model is again prohibitively expensive given the need to deploy a large number of such gateways and can only provide intermittent connectivity to the Internet.

3. MOBILITY TRACES

The mobility model used in this work is based on the actual movement of buses in the King County Metro bus system in Seattle, Washington. a 5100 square kilometer area. The format of traces consists of time, bus id, route id and bus location. The traces capture the bus activities from 8 September 2001 to 9 September 2001. The buses have a highly predictable day-to-day pattern. Here, we first examine some characteristics of bus movement patterns which are independent of system parameters such as the placement of the gateways and the communication range of the wireless devices. Fig. 2 shows that about 90% of the buses have a speed of less than 40 km/h. Since the speed is highly correlated to the degree of mobility, knowledge of bus speeds can be used to predict several important parameters, such as the duration for which routing paths to the mobile gateways will be active, the neighbor list at any node, etc.

Fig. 3 shows the CDF for the distance between neighboring buses. This proves useful for estimating the required radio range and subsequently the wireless technology for inter-bus communication. Fig. 3 indicates that 65% of the buses are less than 1 km away from their nearest neighbor. This suggests that a radio range of 1 km is a judicious choice. Coincidentally, the communication range for DSRC (Dedicated Short Range Communications) [10] is also 1 km.

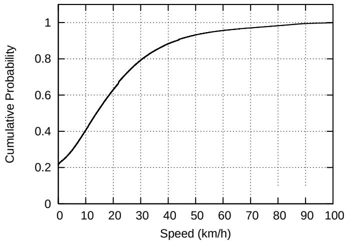

Figure 2: CDF of bus speeds

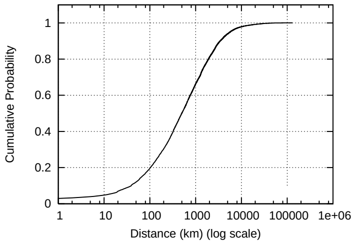

Figure 3: CDF of inter-neighbor distance

4. COVERAGE EVALUATION METRICS

In this section we present the metrics that we use to characterize the connectivity of the buses. Note that, our goal here is to evaluate the extent of coverage possible by deploying mobile gateways on a subset of the entire buses. We are mainly interested in evaluating the reachability at the physical layer independent of the routing protocol in use. Naturally, these characteristics are dependent on the radio range of the wireless devices on each bus. We can broadly classify these metrics into two groups. The first group of metrics (A-D) are used to gain an understanding of the effect of the mobility on various aspects of the inter-bus connectivity such as the link duration and path duration. The second group of metrics (E-F) assume that a certain number of mobile gateways have been deployed on optimally chosen buses, the procedure for which is described below. These metrics then study various aspects of the resulting connectivity that is observed in the network. All of these metrics can provide insights into the expected performance that different applications may perceive once deployed in such a system. Note that in the rest of the paper we will frequently refer to the buses as nodes and the buses equipped with mobile gateways as simply gateways.

4.1 Cluster Size and Number of Clusters

The ad-hoc network that is formed amongst the buses is typically made up of several partitioned network clusters at any given time. Further, the composition of the clusters keeps changing dynamically over time. The cluster size measures the number of buses that constitute each cluster. All nodes within a cluster can reach each of the other nodes via either a single-hop or multi-hop path. Hence, provided the size of a cluster is not very large, it may be sufficient to deploy one gateway for each cluster to ensure basic connectivity to all nodes. For large clusters, it may be necessary to deploy multiple gateways to reduce contention within the cluster. A general rule of thumb suggests that the number of mobile gateways needed would be approximately equal to the number of clusters being formed.

It is also important to get a sense of the number of *orphan clusters*, i,e, the clusters that only have one member. This is because ensuring connectivity to all orphans would warrant installing a gateway atop each of the orphans. Hence, a large percentage of orphan clusters would potentially require the deployment of a large number of mobile gateways.

4.2 Link Duration

This metric indicates the average time that a pair of buses are within each others' radio range. In other words, it shows how long two buses can communicate with each other directly. Since buses are dynamically moving, the links are prone to be broken frequently.

4.3 Path Duration

Path duration refers to the time that one bus can reach another bus and is used to illustrate how long, on average, a path can be maintained. The difference between a path and a link is that a path can be established via either single or multiple hops while a link refers to single-hop communication.

4.4 Longest Path

The longest path denotes the topological distance between two farthest nodes within the same cluster. It gives a sense of the physical dimension of the clusters.

4.5 Percentage of Buses Covered

One would have to deploy at least one mobile gateway in each cluster to provide Internet connectivity to the entire cluster. Note that for very large clusters, multiple gateways would be desired, particularly if the bandwidth required by each node is high. However, in this study we are primarily concerned with providing basic connectivity and hence do not address this issue further. This metric shows the average percentage of nodes that can connect to the Internet, given that a certain number of gateways have been deployed.

Clearly, the choice of the buses on which the mobile gateways are deployed will affect this metric. The buses that are always part of large clusters are ideal gateway candidates. In our evaluations, we have analyzed the clustering patterns of the buses and have devised an optimal gateway placement policy. We first rank the buses according to the average percentage of other nodes that they can provide connectivity to over the entire trace duration. In order to place n gateways in an optimal manner, our algorithm cycles through all n possible combinations from amongst the high ranking nodes to determine the combination that can provide the maximum coverage. Our optimal placement algorithm is particularly feasible for public transport systems, given that the buses always run along fixed routes according to set timetables. The pseudo-codes of our optimal gateway placement algorithm are shown below.

```
FOR i = 0 to totalNumOfBuses - 1
    maxCoverage = 0
    FOR each bus NOT found in SelectedBusForGateway
        SelectedBusForGateway[i] = bus
        value = % of coverage of SelectedBusForGateway
        IF maxCoverage < value THEN
                 maxCoverage = value
                 nextBus = bus
        ENDIF
    ENDFOR
    SelectedBusForGateway[i] = nextBus
ENDFOR
```

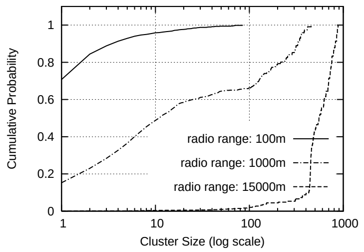

Figure 4: CDF of average cluster size

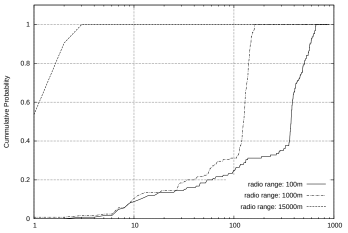

Figure 5: CDF of average number of clusters

4.6 Gateway Path Duration

This metric measures the duration for which an unbroken path exists between a node and a gateway. The average is computed over all paths that existed at least once between a node and a gateway. We use the same optimal gateway placement policy as described above. This metric has a direct implication on the target applications since frequent expiration of the paths will lead to highly intermittent connectivity. However, it can be argued that path breakage may not affect applications if a new one can be found soon enough [7].

5. EVALUATION

In this section we present the results of our analysis of bus traces. Again, our goal here is to evaluate the characteristics of the network connectivity at the physical layer, independent of routing. We also seek to investigate the effect of the radio range on the observed characteristics and hence choose three different ranges: 100m, 1 km and 15 km which correspond to the coverage provided by three candidate access technologies: 802.11, DSRC and 802.16 [11], respectively.

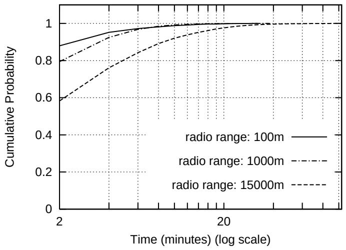

Figure 6: CDF of average link duration

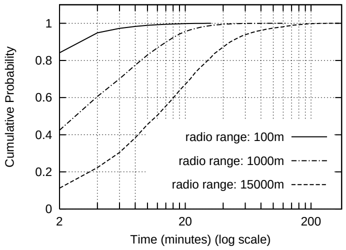

Figure 7: CDF of average path duration

5.1 Cluster Size and Number of Clusters

Fig. 4 shows the CDF for the average cluster size. As seen from the graph, for the 100m radio range, 70% of the buses form orphan clusters and there are only 2 buses in each cluster, on an average. On the other hand, with 1 km, the percentage of orphan clusters significantly reduced to around 15%. Furthermore, there is a significant percentage of clusters with approximately 100 to 400 bus members. These large clusters possibly are formed around city district area at peak times. Overall, the average cluster size is 8 nodes. For the 15 km radio range, as expected, there are no orphan clusters and the cluster size increases dramatically. The ad hoc network is almost always partitioned into 2 or 3 mammoth clusters, with an average of 334 buses per cluster.

5.2 Link Duration

Fig. 6 shows the distribution of the one-hop link durations for different radio ranges. As observed, most of the one-hop links are active for less than 2 minutes. The probability that the link duration is longer than 2 minutes is approximately 13%, 20% and 40% for the 100m, 1km and 15km radio ranges respectively. In general, the link duration increases as the radio range increases. However, the introduction of a longer radio range does not significantly affect the link duration.

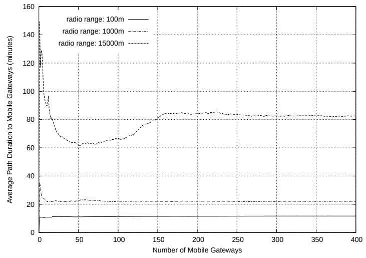

Figure 10: Average gateway path duration

5.3 Path Duration

The CDFs for the path duration are illustrated in Fig. 7. Path duration is always equal to or greater than the link duration since it also accounts for multi-hop paths. For the 100m radio range, approximately 15% of buses have a path duration greater than 2 minutes. However, the path duration increases significantly for the 1km range. The result shows that there is almost a 60% chance that 2 nodes can maintain a path for more than 2 minutes. Finally, as expected, with 15 km the paths remain stable for even longer.

5.4 Longest Path

Fig. 8 shows the average length of the longest path for different cluster sizes for each of the radio ranges under consideration. In the cases of the 100m and 1km radio ranges, we observe that the longest path increases gradually with the cluster size. In addition, the observed longest path is significantly smaller in comparison to the longest path that can be formed amongst the nodes. For example, if 10 nodes are positioned in a chain topology along a straight line, and if the radio range is 100m, then the maximum possible path length is 1000m. However, Fig. 8(a) suggests that on average the longest path for a cluster of size 10 is only around 300m. Fig. 8(a) and 8(b) suggests that this difference is even greater for larger clusters, due to the fact that a large number of the buses are clustered quite close to each other. For the radio range of 15km, most clusters have more than 400 members as shown in Fig. 4. As seen from Fig. 8(c) the average recorded longest path is 70 km across all cluster sizes. Note that, intuitively the longest path in a cluster should increase as the cluster size increase, as in the cases of 100m and 1km radio ranges. However, this is not always true for the 15km radio range. This is because, as observed from Fig. 6, there are many buses with an inter-bus distance of much less than 15 km. Thus, the longest path does not necessarily increase with the cluster size.

5.5 Percentage of Buses Covered

Fig. 9 illustrates the percentage of buses covered as a function of the number of gateways for different radio ranges. We use the optimal gateway placement algorithm as described in Section 4.5 for placing the mobile gateways. As expected for the long range of 15km, only a handful of gateways, 5 to be precise, are sufficient to cover all buses. For the other two radio ranges we observe a gradual increase in the percentage of coverage with an increase in the number of gateways. For example, to achieve 80% coverage we need 470 and 150 gateways for the 100m and 1km radio ranges, respectively. However, for complete coverage the gateways needed increase to 1000 and 800 for these two cases.

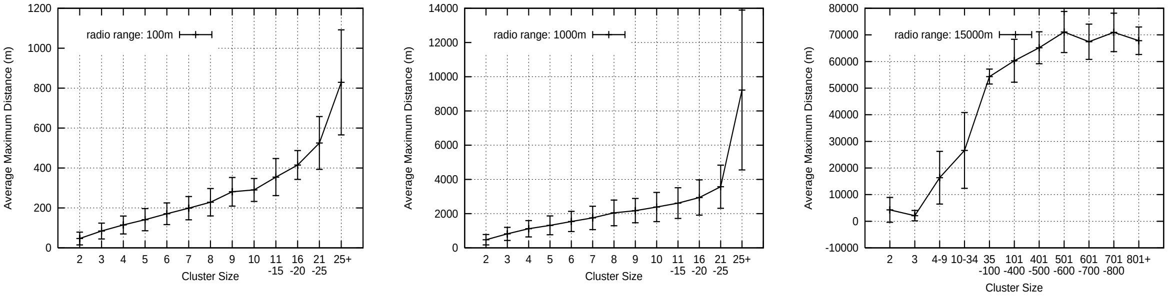

(a) 100m Radio Range (b) 1 km Radio Range (c) 15 km Radio Range

Figure 8: CDF of the longest path

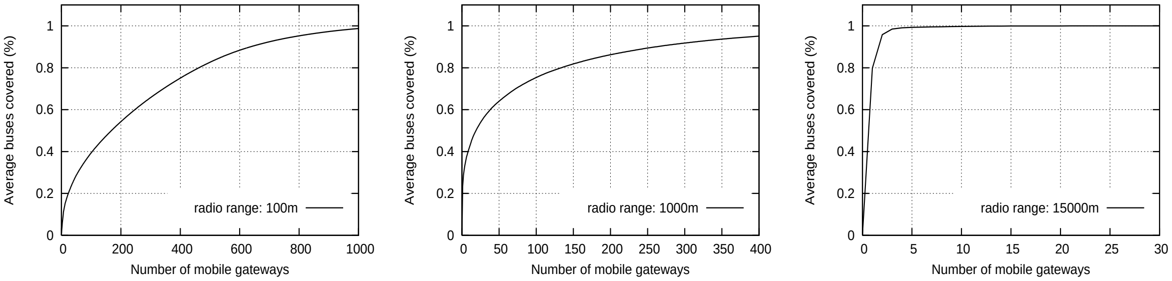

(a) 100m Radio Range (b) 1 km Radio Range (c) 15 km Radio Range

Figure 9: Percentage of buses covered as a function of the number of gateways

5.6 Gateway Path Duration

Fig. 10 indicates that the average duration of a path from a node to a gateway is independent of the number of gateways and is fairly constant. The results are promising for non-real time applications such as web browsing and e-mail which do not require connectivity for a long stretch of time. If a new path to the mobile gateway can be found soon enough following a disconnection, these applications would function in an uninterrupted manner.

6. PREDICTION BASED ROUTING

We now present a novel enhancement based on mobility predictions for improving the performance of multi-hop routing protocols and evaluate its suitability for our architecture. In this discussion, we only focus on the data exchanged between the nodes and the Internet. There have been numerous studies which have evaluated the performance of different routing protocols in forwarding packets between arbitrary source-destination pairs in vehicular ad hoc networks. Since the mobile gateways provide Internet connectivity in our system, data traffic from all the nodes needs to be routed to these gateways. However, as shown in Fig. 10, the node-to-gateway paths break frequently. Hence, to prevent long periods of disconnection, a node needs to quickly establish a new path to another gateway within its current cluster once the old path breaks. In the event of a disconnection, reactive routing protocols such as AODV [13] are usually able to determine that the current path to the gateway is no longer valid after that several retransmission attempts have failed at the link layer. These protocols then try to find a new path to another gateway, provided all gateways are configured to advertise routes to IP addresses external to the bus network. However, these operations do incur significant delays. Moreover, there could be problems if cached routes, which might be no longer valid, were used while establishing routes.

We can leverage the knowledge of the movement patterns of public transport vehicles to aid the routing protocols and overcome the above drawbacks. Since the placement of the gateways is known apriori, a schedule can be constructed off-line for each node, which contains an ordered list of the gateways ranked according to their physical proximity to that node at different time instants. The node can then consult this schedule to determine the destination gateway for the packets at any given time. Of course, in the event that the buses are delayed, the chosen gateway may not be reachable from the node. However, the ranked list allows the node to quickly establish a path to an alternate gateway.

We have carried out some preliminary experiments to test the effectiveness of this predictive policy on several multi-hop routing protocols. The simulations were carried out in ns-2 and we chose to first investigate the performance in a partial region of size 5km x 8km from the large topology covered by the traces. Based on our earlier evaluations in Section 5, we chose a radio range of 1km, consistent with DSRC and use the two-ray ground reflection propagation model. Two gateways were used and their placements were

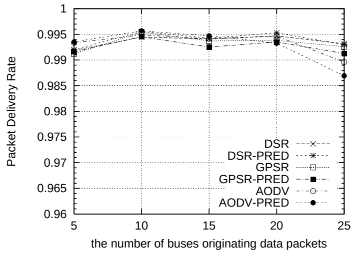

Figure 11: Packet delivery ratio for prediction-enhanced routing

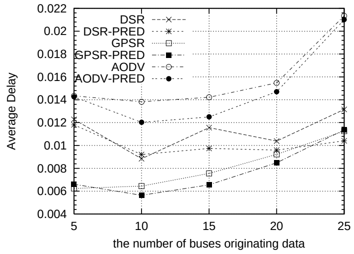

Figure 12: Average Delay for prediction-enhanced routing

determined by the optimal policy discussed earlier. The source nodes were varied from 5 to 30, each of which is a CBR traffic source transmitting packets of size 64 bytes at the rate of 4 packets per second. We used two metrics to evaluate the performance of the routing protocols: (i) packet delivery rate, which is the ratio of the data packets that are successfully delivered to the gateways to the overall number of packets and (ii) average delay, which is measured as the latency for the packet from the source to reach the gateway. The results presented here are averaged over 6 runs, with each simulation lasting for 900 secs. We chose to compare the performance of AODV, DSR [12] and GPSR [14].

Fig. 11 shows that the three routing protocols achieve similar packet delivery rates with DSR exhibiting slightly better performance. Our prediction-based schemes have a similar performance as the original protocols due to the fact that the nodes in our simulations are mostly well-connected. Hence, packet loss is very rare. For a sparser network, we expect a prediction-assisted routing protocol will achieve a higher delivery rate. Fig. 12 shows that our prediction-assisted scheme improves the delay. In our prediction-based scheme, the sources always send packets to the nearest gateway, which results in a lower average delay in comparison to non-prediction based protocols. Since AODV and DSR need to initiate the route discovery before sending data packets, they exhibit a longer delay as compared to GPSR. In addition, DSR uses multi-route caches and eavesdrops route information from data packets. Thus, DSR benefits more from rich cached route information than AODV does and less route discoveries are generated by DSR, which resultingly introduces a shorter delay.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we study the feasibility of providing Internet connectivity for public transport systems using an ad-hoc architecture where packets are routed to the Internet via mobile gateways. For our evaluations, we use realistic mobility patterns of city buses from a metropolitan city. We also analyze the influence of the radio ranges on our results for 3 different wireless standards: 802.11, DSRC and 802.16. Based on our analysis, we conclude that DSRC with a range of 1 km is an ideal candidate for our scenario. With DSRC, the percentage of orphan clusters is less than 20% and most clusters are reasonably small, implying that deploying one gateway per cluster would result in reasonable per-node throughput. By optimally choosing the candidate buses for the gateways, 80% of the buses can connect to the Internet with the deployment of only 150 gateways, about 13% of the overall buses. Further, the average duration of the node to gateway path is mostly greater than 20 minutes, which is sufficient for the smooth functioning of non-real time applications such as Web browsing and e-mail. Finally, we also propose a prediction-based strategy to enhance the performance of routing protocols in our architecture, which uses the known mobility patterns of the buses to predict the closest gateway to any node. Our preliminary results indicate that significant benefits can be reaped by incorporating such mobility-based predictions.

Our results are based on the evaluation of a trace for a single public transport network. An important next step would be to repeat the above analysis for other traces and obtain a set of generalized results. Moreover, the statistical results in this paper are all empirical distributions (i.e. CDFs). A future extension could involve developing analytical models from them so that one can get a better insight into their statistical properties (such as that if some metrics are heavy-tailed).

8. REFERENCES

[1] Icomera, http://www.icomera.com.
[2] PointShot Wireless, http://www.pointshotwireless.com.
[3] NEMO Working Group, http://www.ietf.org/html.charters/nemo-charter.html
[4] V. Namboodiri, M. Agarwal, L. Gao, "A Study on the Feasibility of Mobile Gateways for Vehicular Ad-hoc Networks". In Proceedings of the First ACM International Workshop on Vehicular Ad-Hoc Networks, Philadelphia, June 2004.
[5] E. Huang, W. Hu, J. Crowcroft, I. Wassell, "Towards Commercial Mobile Ad Hoc Network Applications: A Radio Dispatch System", in Proceedings of ACM MOBIHOC, Urbana-Champaign, IL, May 2005.
[6] F. Bai, et al., "IMPORTANT: A Framework to Systematically Analyze the Impact of Mobility on Performance of Routing Protocols for Ad-hoc Networks", in Proceedings of IEEE INFOCOM, March 2003.
[7] J.G. Jetcheva, et al., "Design and Evaluation of a Metropolitan Area Multi-tier Wireless Ad Hoc Network Architecture" , Mobile Computing Systems and Applications, 2003.
[8] J. Ott, D. Kutscher, "Drive-thru Internet: IEEE 802.11b for Automobile Users", in Proceedings of IEEE INFOCOM, Hong Kong, March 2004.
[9] M. Bechler, O. Storz, W. Franz, L. Wolf, "Efficient Discovery of Internet Gateways in Vehicular Communication Systems", in the 57th IEEE Vehicular Technology Conference (VTC), Jeju, South Korea, April 2003.
[10] Dedicated Short Range Communications (DSRC) Service Standard, FCC, http://wireless.fcc.gov/services/its/dsrc/.
[11] E. Crozier and A. Kelin, "WiMAX's Technology for LOS and NLOS Environments", WIMAX Forum.
[12] D. B. Johnson, D. A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks", in Mobile Computing, edited by T. Imielinski and H. Korth, Kluwer Publishing, 1996, pp. 153-181.
[13] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing", IETF RFC 3561, 2003.
[14] B. Karp and H. T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," in Proceedings of MOBICOM 2000, Boston, MA, USA, 2000, pp. 243-254.