

Throughput and Cost-Effectiveness of Vehicular Mesh Networks for Internet Access

Alexandre K. Ligo^{†‡}, Jon M. Peha[†], João Barros[‡]

[†]Carnegie Mellon University
Pittsburgh, USA

[‡]University of Porto
Porto, Portugal

Abstract—Vehicular mesh networks can be used to carry Internet data traffic for mobile users. We use data from a real vehicular network that is operating in Portugal to estimate the relationship between network throughput and offered load of Internet traffic, quantity of vehicles, quantity of infrastructure and the use of Request-to-Send/Clear-to-Send (RTS/CTS) handshaking. We show that achievable throughput remains close to its maximum even for high levels of offered load per vehicle or high density of vehicles, so congestion control mechanisms in mesh networks can effectively prevent throughput from collapsing as these factors increase over time. This achievable throughput can be increased with the deployment of additional roadside infrastructure that serves as a gateway to the Internet, although the throughput gain per gateway decreases as more are added. Deploying these gateways is cost-effective, i.e. economic benefits of the resulting throughput exceed costs of infrastructure, when vehicle density is sufficiently high. We also find that use of RTS/CTS decreases achievable throughput in this scenario.

Keywords—vehicular networks; DSRC; mobile Internet; network throughput; simulation; cost-effectiveness

I. INTRODUCTION

Internet traffic has been growing sharply in mobile networks [1], and vehicular mesh networks can be used as a new and cost-effective way to provide Internet access. Such networks may soon be widely deployed using Dedicated Short Range Communications (DSRC) technology, primarily as a way of enhancing automotive safety. Routers in a vehicular network are placed in automobiles and in infrastructure placed near roads, and thus the latter can provide Internet access for users in vehicles when the DSRC channels are not needed for safety-related applications, as illustrated in Fig. 1. This paper examines throughput of Internet access via DSRC-based vehicular networks, and economic benefits and costs of such a use as well.

Using packet-level simulation in combination with data collected from an actual city-wide vehicular network that is operating in Porto, Portugal, we address two interrelated research questions. The first is how throughput of vehicular mesh networks between DSRC-equipped vehicles and roadside infrastructure connected to the Internet depends on factors such as the quantity of vehicles, the quantity of infrastructure, offered load, and whether Request-to-Send and Clear-to-Send (RTS/CTS) handshaking is used. It is often difficult to predict the achievable throughput of a mesh network, even if nodes are stationary as with some citywide networks, and these nodes are moving. We expect the number of DSRC-equipped vehicles to

increase over time, as will the load of Internet traffic per vehicle, which means there is a possibility of ever-increasing interference and congestion that reduces the achievable throughput. A variety of mechanisms including RTS/CTS and exponential back-off are intended to limit the damage of congestion and interference. On the other hand, unlike cellular networks, mesh networks such as these can exhibit *cooperative gain* [2]–[5], whereby capacity actually increases as more devices are added. By analyzing how various factors affect achievable throughput, this paper will shed light on how trends like these might affect performance in mesh networks that are connected to infrastructure of gateways to the Internet, which is the case for vehicular networks connected to roadside infrastructure, and other types of networks as well such as mesh-based metropolitan-area networks.

Of course, performance is not the only consideration for such decisions. This brings us to the second research question, which is when deployment of more roadside infrastructure is cost-effective, i.e. when the incremental benefit of adding roadside units (RSUs) exceeds the incremental cost. For this analysis, we assume that each Mbps of throughput that can be sustained over the vehicular network has a given value, and that each RSU has a given cost, regardless of where devices are located. In many cases, this value per Mbps of increasing capacity of a vehicular network would be the cost of carrying that same traffic over a cellular network instead [6]. We do not assume that a vehicular network must provide ubiquitous coverage at all times, so devices that are temporarily disconnected would either wait until they are connected to transmit or switch to a ubiquitous cellular service. In contrast, previous research [7] has considered the question of whether a vehicular network is less expensive than a cellular network when both are expected to be equally ubiquitous, and when neither the vehicular nor the cellular network is ubiquitous. They find that vehicular networks are cost-effective for the specific set of assumptions that they considered.

While previous research has used theoretical analysis based on abstract models (e.g. [7]), we use simulation with assumptions derived from empirical measurements to employ realistic representations of the elements of a vehicular network that greatly affect throughput, including the location of



Fig. 1. Vehicular mesh network with roadside infrastructure and multihop connections for Internet access.

This work is supported in part by the CMU-Portugal Partnership (scholarship SFRH/BD/51564/2011) and Portugal Foundation for the Science and Technology (ref. UID/EEA/50008/2013), in part by Carnegie Mellon University's Technologies for Safe and Efficient Transportation, the National USDOT University Transportation Center for Safety (T-SET UTC) sponsored by the US Department of Transportation, and in part by the National Science Foundation under Grant No.1547237.

roadside units (RSUs) and vehicles equipped with onboard units (OBUs), the signal loss, and the DSRC protocol itself. The simulation results show the relationship between throughput of the vehicular network and the quantity of OBUs, RSUs, offered load and the use of RTS/CTS.

This paper is organized as follows. Section II describes the Porto dataset, while Section III describes the methodology and results of the throughput analysis. In Section IV we discuss economic benefits and costs. Conclusions and final remarks are described in Section V.

II. DATASET

We use data from a real DSRC network that is operating in the city of Porto, in Portugal. As of March 2015, 430+ buses have offered free Wi-Fi to passengers, and the Wi-Fi traffic generated in each bus is routed through a DSRC OBU, which is sent from vehicle to vehicle until data reach one of 27 DSRC RSUs installed in fixed locations such as traffic lights. Data is sent over cellular only when a DSRC connection is not available. In this network, passengers have been transferring about 3 TB per month, and up to 70% of data is carried via the vehicular network, especially in downtown where vehicles are concentrated and most of the RSUs are placed. We used a dataset with measurements of data transferred over DSRC and cellular, and also position data of the 850+ buses and taxis over a 41 km² area. Porto data is used in two ways: first, bus and taxi GPS positions from Porto are used to determine the positions of the vehicles in the simulation every 5 seconds. Second, the received signal measured in the buses is used to verify whether the simulated signal loss is compatible with measured loss, on average.

III. THROUGHPUT ANALYSIS

In this section we examine the relationship between the throughput of DSRC-based vehicular networks and offered load, density of OBUs and RSUs, and the use of RTS/CTS.

A. Methodology for Throughput Estimation

We use the aforementioned location data of buses and taxis to simulate a mesh network comprised by OBU-equipped vehicles that connect to RSUs to exchange Internet traffic. We use this simulation to estimate the maximum rate of data that can be carried through DSRC under several conditions.

Transfers of data packets are simulated between pairs of vehicles and between vehicles and RSUs at the transport, network, link and physical layers using the ns-3 network simulator [8]. At any given time, packet streams flow between each vehicle and one RSU which serves as a gateway to the Internet, either directly or through multiple hops with other vehicles acting as relays.

Nodes in the simulation represent DSRC-equipped buses, taxis, passenger cars, and RSUs. The positions of buses and taxis every 5 seconds are derived from the GPS logs from Porto buses and taxis. We assume that at peak hour, 1.3% of vehicles on the road are buses, as derived from [9], [10]. Each simulated bus follows the same trajectory as a real bus. We randomly select this bus and a start date and time, and then use its actual GPS measurements. In our simulation, each of the remaining vehicles follows the same trajectory as a real taxi.

For these, we similarly select a random taxi, and a random start day and time.

RSUs are placed in locations that are close to a large number of vehicles at peak hours, using the *k-means clustering heuristic* [11]. This popular algorithm first divides a number of observations (vehicle locations, in our case) into *k* regions, and then finds the centroid for each region that minimizes the sum of the distances between the observations (vehicles) and the centroid. RSUs are placed at these centroid positions.

Bus antennas have a height of 3 meters (average of single deck buses in Porto), and all other vehicles have height of 1.5 meters (which is consistent with [12]). RSU antennas have a height of 7 meters (average in Porto).

The network is simulated as representing a network with non-moving nodes communicating for 5-second intervals, after which vehicle positions are changed and the network is simulated for the next interval.

Each RSU is a gateway to the Internet which a given vehicle connects to, but we treat the RSU as if it were the endpoint of a TCP connection rather than merely a gateway. The throughput rate via DSRC is defined as the sum, across all OBU-equipped vehicles, of the data throughput achievable between each vehicle and an RSU it is connected to (either through a single or multiple hops), with all such flows at steady-state throughput in each 5-second interval. For each vehicle, 90% of all data flow downstream and 10% upstream.

The DSRC band in the U.S. is divided in seven 10 MHz channels (172 to 184) of which channels 172 and 184 are reserved for safety applications, and channel 178 is allocated for control of operation of the other channels [13]. The remaining four channels are thus assumed as available for Internet access, with each vehicle OBU and each RSU being equipped with four radios.

The assumed TCP Maximum Segment Size is 2244 bytes (based on typical TCP connections traversing 802.11). IP packets are routed through the path with the minimum number of hops between the vehicle and a RSU, up to a maximum of three hops for each path. If a given vehicle can reach one or more RSUs through one-hop paths, then the path with the least path loss is selected. If the minimum number of hops in all paths is greater than one, then we select one path randomly among the paths with the minimum number of hops, such that each of those paths is equally likely. This method is a simplification in that it assumes that every vehicle learns immediately about all paths with three hops or less through some unspecified routing protocol. (There is ongoing research on routing protocols for these networks [14]–[17].)

When a vehicle OBU is not within three hops of a RSU, the OBU will either wait until it is connected to a RSU, or switch to a ubiquitous cellular service. Thus, it is assumed that either applications are tolerant to the periods of disconnection, or disruptions from switching to/from cellular are negligible. (DSRC links can be established in just 300 milliseconds [18], [19]).

The media access control (MAC) sublayer in DSRC is the one specified in IEEE 802.11p [18], with all packets transferred

having the same priority with respect to the IEEE 802.11p user priority levels. RTS/CTS handshaking is not used, unless otherwise stated.

The assumptions for the physical layer are as follows. The receiver sensitivity threshold is -94 dBm. A link is established between two nodes if and only if signal strength at the receiver exceeds 15 dB above the sensitivity threshold. (This is the criteria determined empirically in the Porto network.) Packets are received at an error rate ([8],[20]) that depends on the signal-to-interference-plus-noise ratio (SINR), which determines the modulation to be used, allowing bitrates of up to 27 Mbps per channel in IEEE 802.11p. The transmit power is 14.6 dBm, obtained from measurements at the equipment output (which is also consistent with [21], [22]), and the gains of the transmission antennas are 16 dBi and 5 dBi for the RSUs and vehicles, respectively (consistent with the equipment in the Porto network). The received signal is calculated according to the propagation loss model from [23] (urban microcell B1 variant). This model includes parameters that are relevant for vehicular networks, such as the height of vehicle and RSU antennas. Moreover, it was designed for and is most accurate in urban areas. This is appropriate because demand for Internet capacity obtained via DSRC is greater in urban areas, since mobile Internet is currently provided through macrocellular and because vehicular networks are more likely to be deployed in urban areas where there are more cars [6].

The difference between the median simulated loss and the median loss measured in Porto buses is less than 5 dB for most distances shorter than 200 meters, which shows the assumed model and parameters provide a reasonable approximation of observed loss. For example, at a distance of 100 m between a RSU and a bus, the median measured loss is 92 dB while the simulated loss is 95 dB. More than 95% of the hops observed in the Porto network are shorter than 200 meters.

B. Throughput of DSRC Vehicular Networks in Urban Areas

The simulation allows us to examine how the throughput of a vehicular network is affected by several factors, including offered load, the quantities of OBUs and RSUs, and the use of RTS/CTS. For the graphs in this section, the assumptions are those described in Section III.A, unless otherwise stated. Each data point is derived by averaging throughput for 1000 vehicles in a simulated 20 km² region over a 50-second period. If we

make the simplifying assumption that the throughputs of these 1000 vehicles are mutually independent, then the confidence interval is within 7 to 15% of the mean throughput across all vehicles.

Fig. 2 shows the relationship between throughput per unit of area and offered load per OBU sent over the vehicular network. Curves for several RSU densities are shown, and the density of OBU-equipped vehicles is fixed at 100 OBUs/km² for all curves. The graph shows that as load increases, throughput increases rapidly until it reaches a peak, and then throughput remains within a small percentage of its peak for higher loads, regardless of RSU density. At low loads, there is little interference or congestion, so throughput increases linearly in this range. More importantly, the fact that throughput is close to peak even for much higher loads means that congestion and interference never cause a serious loss of throughput, probably thanks to mechanisms such as MAC-level listen-before talk and exponential back-off of the Distributed Coordination Function (DCF) of IEEE 802.11 [24], and transport-layer dynamic windows. As a result, cities with vehicular networks need not fear that achievable throughput will get worse as load per car goes up every year, as it is expected to do [1].

Researchers have previously explored the relationship between offered load per device and throughput in the much simpler cases of a single wireless local-area network using Wi-Fi [24], [25], as opposed to a vast mesh network like ours. They similarly observed that throughput first increases with offered load per device, and then plateaus at a level close to the peak. They concluded that this *saturation throughput*, which is defined as throughput approached under high offered load, depends on MAC-level DCF, which is used both in Wi-Fi and DSRC [18]. That is probably a factor in vehicular mesh networks as well.

Fig. 3 shows saturation throughput per unit of area as a function of the density of vehicles for different RSU densities. Fig. 3 shows that saturation throughput increases linearly when OBU density is low, and then remains close to its maximum for all OBU densities above some threshold. This is important, because the number of DSRC-equipped vehicles will increase over time if the U.S. Dept. of Transportation mandates the technology for all new cars [26].




Fig. 2. Throughput of the vehicular network for varying load of Internet data and 100 OBUs/km². Each line refers to a fixed RSU density.




Fig. 3. Saturation throughput of the vehicular network for varying density of OBU-equipped vehicles. Each line refers to a fixed RSU density.

The curves in Fig. 3 look similar to Fig. 2, but changing the number of DSRC-equipped vehicles per area has an impact beyond just increasing offered load per area. It could also create more opportunities for devices to interfere with their neighbors in a way that decreases throughput per area, as has been observed within a single wireless LAN [24], or more opportunities for devices to *cooperate* with each other in a way that increases throughput per area, which is known as cooperative gain. For example, it has been shown [5] that aggregate capacity of an ad hoc wireless network can increase indefinitely with the square root of the number of nodes, when each pair of nodes within the ad hoc network is communicating with each other (with aggregate capacity defined as the sum across all possible node pairs of achievable throughput times distance). However, we do not observe this same cooperative gain in Fig. 3, probably because of the following differences. First, in our work throughput in a vehicular network used for Internet access is measured as throughput between the mesh and the Internet, rather than throughput within the mesh network, and the number of RSUs acting as gateways between mesh and Internet is held constant. Second, we assume a common power level across all transmissions, whereas in [5] each vehicle may choose a different level. Third, we assume a realistic geographic distribution of vehicles, whereas in [5] vehicles are uniformly distributed.

As Fig. 2 and Fig. 3 show, throughput per km^2 always increases as RSUs are added, for all levels of load and for all OBU densities examined. A new RSU can both increase coverage and reach vehicles that were previously disconnected, and it can provide shorter and therefore higher-capacity links (because of improved SINR) to vehicles that were already connected. Thus, as offered load per vehicle and penetration of OBUs in vehicles grow over time, more RSUs can be deployed to carry additional Internet traffic.

However, some RSUs have greater impact on saturation throughput than others. This is shown in Fig. 4, which shows average saturation throughput per RSU as a function of RSU density, for two different OBU densities. The graph shows that when RSU density is low, the throughput per RSU is high. For example, the theoretical maximum link bitrate for an RSU that is situated very close to all the devices it serves is 108 Mbps, and we see average throughputs greater than a third of this




Fig. 4. Average saturation throughput per RSU, for varying density of RSU infrastructure. Each line refers to fixed OBU density.

unachievable limit. However, much higher RSU densities yield lower throughput per RSU, even though more RSUs results in shorter distances to nearby vehicles.

A major reason for this is the k-means heuristic used to select the location of RSUs, which attempts to place each new RSU in a place that will support high throughput. As more RSUs are placed, the algorithm must choose less desirable locations. Thus, the fact that throughput per RSU decreases rapidly for RSU densities under 3 per km^2 is one indicator that the k-means placement algorithm is working well. At RSU densities around 6 per km^2 , the majority of vehicles are served by a single-hop path to an RSU, so additional RSUs have only modest benefit, and throughput per RSU would fall for almost any RSU placement algorithm. This is consistent with previous studies of the effect of infrastructure on the throughput of vehicular networks such as [27]. However, the aforementioned work is based on a simple graph-theoretical model with perfect, collision-free scheduling and link rates determined by distance only. In contrast, our model considers DCF scheduling with its associated delays, and link rates and transmission errors based on an empirical urban channel model [23]. These differences make our model more representative of a DSRC network in an urban environment.

Throughput per RSU decreases as more RSUs are deployed, but the cost of each is constant. Thus, we eventually reach a point where it is no longer cost-effective to deploy another RSU. This issue will be explored further in Section IV.

Protocol mechanisms such as DCF do minimize packet collisions, but they do not avoid the hidden node problem that can be present in vehicular networks. This problem can be overcome by RTS/CTS handshaking [24], which minimizes the probability of collisions caused by hidden nodes. Fig. 5 shows saturation throughput per unit of area as a function of RSU density, with one curve showing results when RTS/CTS is used for all transmissions, while the other curve shows results without RTS/CTS.

Fig. 5 shows that the use of RTS/CTS handshaking greatly decreases saturation throughput in this scenario. While this result is consistent with previous work such as [28], it




Fig. 5. Saturation throughput for varying density of RSU infrastructure, and 50 OBU/km². One line refers to basic access (no RTS/CTS handshaking), and the other refers to RTS/CTS used for all transmissions.

contradicts others such as [24]. One reason is that RTS and CTS frames are sent at the minimum link bitrate, which slows transmissions of data frames that are sent at link bitrates close to the theoretical maximum [28]. Another reason is that each RTS and CTS frame introduces a fixed idle time to the protocol. Both reasons reduce throughput, and that reduction may exceed any gains in avoiding collisions from hidden nodes, especially when the probability of such collisions is low. Designers of vehicular networks for Internet access should not use RTS/CTS to maximize throughput, under any level of RSU deployment, as long as the network conditions are similar to the scenario simulated in this paper. Use of RTS/CTS under different conditions, or combined with other techniques to improve its efficiency such as packet concatenation (similarly as in [29]), are subjects for future work.

The results discussed are based on assumptions that reflect vehicular DSRC mesh networks with Internet-connected RSUs, but similar results are likely in other networks with the same key characteristics. For the purpose of throughput analysis, the networks that we considered are similar to mesh networks such as metropolitan-area networks with gateways to the Internet. Moreover, the MAC-level mechanisms to handle interference and congestion discussed in this paper are common to any IEEE 802.11 networks, while TCP congestion control is present in most Internet end-to-end connections. Therefore, although the numerical results depend on the specific network characteristics, it is likely that other mesh networks can also provide a saturation throughput that is close to a peak when the number of devices and the offered load per device grow to high values, and that the saturation throughput can be increased by adding more Internet-connected gateways.

IV. BENEFITS AND COSTS

In this section, we investigate whether the Internet throughputs delivered through vehicular networks are enough to be cost-effective.

A. Assumptions for Benefits and Costs

The analysis of benefits and costs is performed according to the approach and assumptions adopted in [6]. We consider the case where spectrum has already been allocated for vehicular networks, as is the case in the U.S. and many other regions




Fig. 6. Present values of benefit and RSU cost for varying OBU density, and load of 400 kbps per OBU-equipped vehicle on the road at peak hours.

worldwide, and where there is a mandate to equip new cars with OBUs for safety purposes, as is likely to occur soon in the U.S. These resources can be leveraged to provide Internet access via vehicular networks by deploying RSUs that can serve as Internet gateways. The question then is whether the benefits of Internet access justify the costs.

A single RSU can be deployed at a one-time cost of \$1,200-4,000 [6], [30], but there are many costs beyond this initial deployment such as ongoing expenses for backhaul, energy and maintenance. We assume a total lifecycle cost of \$14,000 per RSU, which corresponds to an estimate [6] of the net-present value (NPV) of RSU costs over a ten-year period.

Economic benefits are proportional to the throughputs that can be sustained, which in the previous Section were shown to depend on offered load. We assume that in any 5-second interval during the peak hour, 50% of the DSRC-equipped vehicles on the road are endpoints for data sent continually at 800 kbps (total downstream and upstream). The remaining vehicles are not endpoints for traffic, although they may relay packets for other vehicles in multihop connections. Thus the average traffic per DSRC-equipped vehicle on the road is 400 kbps. This is consistent with the Deutsche Telekom prediction that vehicular traffic will reach 5 GB/month in some years [1], if the average vehicle is on the road 57 min per day [9], and average data rate is the same whenever the vehicle is in use. In reality, data rates vary from vehicle to vehicle, but since RSUs are typically in range of dozens of DSRC-equipped vehicles at all times during peak hour, this simplifying assumption should have limited effect on aggregate throughput.

We also assume that the capability to support t Mbps throughout the peak hour every single day is worth $t * \$2600$, which corresponds to an estimate [6] of the NPV of what it would cost over a ten-year period to add this capability to a typical cellular network.

B. Analysis of Benefits and Costs of DSRC Vehicular Networks

The benefit of a vehicular network will depend greatly on the density of DSRC-equipped vehicles on the road. Fig. 6 shows the NPV of both benefit and cost per square km as a function of density of OBU-equipped vehicles. Other assumptions in Fig. 6 (and subsequent figures) are those




Fig. 7. Present values of benefit and RSU cost for varying load per OBU-equipped vehicle, and 50 OBU-equipped vehicles on the road at peak hours.

described in Sections III.A and IV.A, unless otherwise stated. For the values of OBU density in which the NPV of benefit of Internet access exceeds the NPV of cost of RSUs, the density of RSUs chosen is the quantity that maximizes the difference between the NPV of benefit and the NPV of RSU cost. For the OBU density values in which the NPV of benefit is lower than the NPV of RSU cost for any quantity of RSUs, the optimal density is obviously zero. However, for these points we instead show the density of RSUs as a linear extrapolation from the densities of RSUs calculated over the OBU density range in which the NPV of benefit is greater than the NPV of RSU cost.

Fig. 6 shows that benefit increases faster than RSU cost, and the former eventually exceeds the latter after a sufficiently high OBU density. The reason is that for the assumed load, throughput grows roughly proportionally to OBU density but the optimal number of RSUs rises at a lower pace (because having more RSUs does not necessarily result in a proportional increase in throughput, as shown in Section III). If OBU density is correlated with population at a given time, then infrastructure for Internet access would be cost-effective, i.e. yield a benefit that exceeds RSU costs, in more densely populated areas. Fig. 5 shows that this would happen where there are more than 40 OBUs per km², which in [6] we estimate as the average at peak hours in densely populated cities such as Chicago and Porto for OBU penetration of 25%, which is what we might expect a few years into an OBU mandate for safety. Moreover, OBU penetration in vehicles is expected to increase over time. Therefore, there will be locations at which OBU density is not enough to make vehicular networks cost-effective for Internet access in the early years, but deployment will become cost-effective after a few more years.

The other factor that is likely to change rapidly over time is offered load. Fig. 7 shows the NPV of benefit and RSU cost as a function of the load of Internet data per OBU-equipped vehicle. The density of RSUs for each offered load is chosen as in the previous graph. Fig. 7 shows that the difference between benefit and RSU cost increases with load of traffic per vehicle. If load is going to increase over time as some predict, there will be locations at which OBU density is initially not enough to result in benefit that is higher than cost, but where RSU deployment becomes cost-effective as soon as offered load grows. For a sufficiently high data rate, benefit exceeds cost.

For a more detailed examination of whether and when vehicular networks would be a cost-effective alternative for providing Internet access when compared to cellular networks, see [6]. That paper includes further analysis of how these and other factors affect the value of carrying Internet traffic over a vehicular network rather than over a macrocellular network. To help inform decisions on spectrum allocation and on whether to mandate OBUs, the paper also investigates whether these benefits exceed other DSRC-related costs, such as the cost of OBUs and the cost of spectrum.

V. CONCLUSION

In this paper we use data from an actual city-wide vehicular network that is used to provide Internet access in Portugal, and examine the relationship between throughput of vehicular mesh networks and the load of Internet traffic, the quantity of OBU-

equipped vehicles, the quantity of Internet-connected RSUs, and the use of RTS/CTS handshaking. Moreover, we examine benefits and costs to deploy RSU infrastructure for given quantities of OBU-equipped vehicles and load of Internet traffic.

When offered load per vehicle is small, throughput increases with load. When load is high, throughput approaches a saturation level that remains close to the maximum achievable throughput both for high levels of load and for high OBU densities. This implies that protocol mechanisms such as DCF collision avoidance and TCP congestion control will help vehicular networks to sustain stable throughput, even as the load of Internet traffic and the penetration of OBUs in vehicles grow sharply as predicted.

We also show that deploying more RSUs increases throughput for all levels of load and OBU density examined, although the incremental gain decreases with additional RSUs. This happens because RSUs are placed first in locations where they can provide the highest throughput, and then additional RSUs yield lower throughput than the first ones. Eventually, adding extra RSUs results in little incremental throughput. One implication is that if the cost per RSU is constant, there may be a number of RSUs above which further deployment is not worth the lower incremental throughput.

Another factor that affects the throughput of DSRC and other 802.11 networks is the possibility of collisions caused by hidden nodes, which can be reduced by the use of RTS/CTS. However, the mechanism itself introduces delays that may offset the gain from avoided collisions. We show that RTS/CTS actually reduces saturation throughput significantly in the scenarios considered in this paper.

Although some input parameters were selected specifically to be realistic in vehicular networks, the general results above are likely to be applicable in other mesh networks that are connected to an infrastructure of gateways to the Internet, as has been proposed for mesh-based metropolitan-area networks. In particular, it is likely that these mesh networks can also provide a saturation throughput that is close to the maximum even when the number of devices in the mesh and the offered load per device grow large, and that this saturation throughput can be increased by adding more gateways to the Internet.

We also estimate that economic benefits of Internet access through vehicular DSRC networks will exceed the cost of RSU deployment for OBU densities that previous work [6] estimates will be reached in some big cities within just a few years after an OBU mandate [31]. Moreover, as load of Internet traffic and OBU penetration increase over time, benefits will exceed costs for ever less populated areas.

ACKNOWLEDGMENT

The authors thank Veniam Networks, Instituto de Telecomunicações – Porto, STCP, the municipality of Porto, and the Future Cities Project (European Commission EU FP7 under grant number 316296), for the data and support provided.

REFERENCES

- [1] Deutsche Telekom, "Connected cars get big data rolling," 2013. [Online]. Available: <http://www.telekom.com/media/media-kits/179806>. [Accessed: 03-Apr-2015].
- [2] J. M. Peha, "Sharing spectrum through spectrum policy reform and cognitive radio," *Proceedings of the IEEE*, 2009. [Online]. Available: www.ece.cmu.edu/~peha/papers.html.
- [3] H. J. Kim and J. M. Peha, "Detecting Selfish Behavior in a Cooperative Commons," *3rd IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks DYSpan*, 2008. [Online]. Available: www.ece.cmu.edu/~peha/papers.html.
- [4] A. Agarwal and P. R. Kumar, "Capacity bounds for ad hoc and hybrid wireless networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 3, p. 71, 2004.
- [5] F. Xue and P. R. Kumar, "Scaling Laws for Ad Hoc Wireless Networks: An Information Theoretic Approach," *Found. Trends Netw.*, vol. 1, no. 2, pp. 145–270, 2006.
- [6] A. K. Ligo, J. M. Peha, P. Ferreira, and J. Barros, "Comparison between Benefits and Costs of Offload of Mobile Internet Traffic Via Vehicular Networks," *43rd Telecommunications Policy Research Conference (TPRC)*, 2015. [Online]. Available: www.ece.cmu.edu/~peha/papers.html.
- [7] N. Lu, N. Zhang, N. Cheng, X. Shen, J. W. Mark, and F. Bai, "Vehicles Meet Infrastructure: Toward Capacity–Cost Tradeoffs for Vehicular Access Networks," *IEEE Trans. Intell. Transp. Syst.*, vol. 14, no. 3, pp. 1266–1277, 2013.
- [8] "ns-3 Network Simulator." [Online]. Available: <https://www.nsnam.org/>. [Accessed: 10-Jul-2015].
- [9] A. Santos, N. McGuckin, H. Y. Nakamoto, D. Gray, and S. Liss, "Summary of Travel Trends: 2009 National Household Travel Survey," 2011.
- [10] European Commission, *EU transport in figures - Statistical pocketbook 2012*, 2012.
- [11] A. W. Moore, "K-means and Hierarchical Clustering," 2001. [Online]. Available: <http://www-2.cs.cmu.edu/~awm/tutorials/kmeans.html>. [Accessed: 30-Jul-2015].
- [12] M. Boban, T. T. V. Vinhoza, M. Ferreira, J. Barros, and O. K. Tonguz, "Impact of Vehicles as Obstacles in Vehicular Ad Hoc Networks," *IEEE J. Sel. Areas Commun.*, vol. 29, pp. 15–28, Jan. 2011.
- [13] C. Campolo and A. Molinaro, "Multichannel communications in vehicular Ad Hoc networks: a survey," *IEEE Commun. Mag.*, vol. 51, no. 5, pp. 158–169, May 2013.
- [14] R. Meireles, "Leveraging Diversity and Spatial Connectivity in Multi-hop Vehicular Networks," Carnegie Mellon University, University of Porto, 2015.
- [15] F. Li and Y. Wang, "Routing in vehicular ad hoc networks: A survey," *IEEE Vehicular Technology Magazine*, vol. 2, pp. 12–22, 2007.
- [16] R. Meireles, P. Steenkiste, and J. Barros, "DAZL: Density-aware zone-based packet forwarding in vehicular networks," in *IEEE Vehicular Networking Conference, VNC*, 2012, pp. 234–241.
- [17] N. Wisitpongphan, F. Bai, P. Mudalige, V. Sadekar, and O. K. Tonguz, "Routing in Sparse Vehicular Ad Hoc Wireless Networks," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 8, Oct. 2007.
- [18] IEEE, "802.11p-2010 - IEEE Standard for Information technology-- Local and metropolitan area networks-- Specific requirements-- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular En." IEEE, 2010.
- [19] Q. B. Mussabbir and W. Yao, "Optimized FMIPv6 Using IEEE 802.21 MIH Services in Vehicular Networks," *IEEE Trans. Veh. Technol.*, vol. 56, no. 6, pp. 3397–3407, 2007.
- [20] M. Lacage and T. R. Henderson, "Yet another network simulator," in *Proceeding from the 2006 workshop on ns-2: the IP network simulator - WNS2 '06*, 2006, p. 12.
- [21] A. Cardote, F. Neves, S. Sargento, and P. Steenkiste, "A statistical channel model for realistic simulation in VANET," in *2012 IEEE Vehicular Networking Conference (VNC)*, 2012, pp. 48–55.
- [22] F. Bai, D. D. Stancil, and H. Krishnan, "Toward understanding characteristics of dedicated short range communications (DSRC) from a perspective of vehicular network engineers," in *Proceedings of the sixteenth annual international conference on Mobile computing and networking - MobiCom '10*, 2010, p. 329.
- [23] J. Meinilä, P. Kyösti, T. Jämsä, and L. Hentilä, "WINNER II Channel Models," in *Radio Technologies and Concepts for IMT-Advanced*, M. Döttling, W. Mohr, and A. Osseiran, Eds. Wiley, 2009.
- [24] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *JSAC, IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, 2000.
- [25] H. Chhaya and S. Gupta, "Performance modeling of asynchronous data transfer methods of IEEE 802.11 MAC protocol," *Wirel. networks*, vol. 3, pp. 217–234, 1997.
- [26] U.S. Department of Transportation, "Planning for the Future of Transportation: Connected Vehicles and ITS," 2015.
- [27] F. Malandrino, C. Casetti, C.-F. Chiasserini, and M. Fiore, "Content Downloading in Vehicular Networks: What Really Matters," in *IEEE INFOCOM*, 2011.
- [28] P. Chatzimisios, A. C. Boucouvalas, and V. Vitsas, "Effectiveness of RTS/CTS handshake in IEEE 802.11a Wireless LANs," *Electron. Lett.*, vol. 40, no. 14, pp. 915–916, Jul. 2004.
- [29] N. Barroca, M. Borges, and F. J. Velez, "IEEE 802.15.4 MAC Layer Performance Enhancement by employing RTS/CTS combined with Packet Concatenation," in *IEEE ICC - Ad-hoc and Sensor Networking Symposium*, 2014, pp. 466–471.
- [30] J. Wright, J. K. Garrett, C. J. Hill, G. D. Krueger, J. H. Evans, S. Andrews, C. K. Wilson, R. Rajbhandari, and B. Burkhard, "National Connected Vehicle Field Infrastructure Footprint Analysis," US DOT - FHWA, 2014.
- [31] J. Harding, G. Powell, R. Yoon, J. Fikentscher, C. Doyle, D. Sade, M. Lukuc, J. Simons, and J. Wang, "Vehicle-to-Vehicle Communications: Readiness of V2V Technology for Application," Washington D.C., 2014.