Spectrum for V2X: Allocation and Sharing

Alexandre K. Ligo, *Member, IEEE*, and Jon M. Peha, *Fellow, IEEE*

***Abstract*—This paper investigates how much spectrum should be available for intelligent transportation systems (ITS), and whether part of that spectrum should be shared with unlicensed devices, as has been considered by the U.S. Federal Communications Commission (FCC), and if so, what sharing scheme should be adopted. We found that the ITS bandwidth that maximizes social welfare could be either much more or much less than what has already been allocated, because optimal bandwidth is sensitive to uncertain factors, such as device penetration, future data rates, and spectrum opportunity cost. That uncertainty is offset if ITS spectrum is shared under a scheme of coexistence among equals. We also found that the bandwidth required to obtain given throughputs on shared spectrum can be considerably less than the bandwidth to obtain the same throughputs in separate bands. We conclude that the spectrum available for ITS should be maintained or increased, but much of ITS spectrum should be shared with non-ITS devices.**

***Index Terms*—Spectrum sharing, unlicensed spectrum, ITS, connected vehicles, DSRC, V2X, policy.**

I. INTRODUCTION

THE FCC has allocated 75 MHz of spectrum in the 5.9 GHz band for ITS [1], and has adopted the Dedicated Short Range Communications (DSRC) standards for vehicle-to-everything (V2X) communications [2]. V2X can support road safety and non-safety applications [3]. Yet, the question of whether ITS should have an exclusive allocation of 75 MHz is hotly debated. One issue that we address in this paper is how much spectrum should be made available for ITS, rather than other purposes. A related issue that we address is whether the ITS band should be shared with non-vehicular devices, and if so, how. The FCC issued a Notice of Proposed Rulemaking (NPRM) to permit unlicensed devices in that band [4]. However, to date there has been no consensus on the rules to be adopted for such sharing [1].

We address the spectrum issues above by looking into several interrelated research questions. On the issue of how much spectrum to allocate for ITS, we assume that a certain amount of spectrum is sufficient to serve road safety applications, and then explore whether adding spectrum would result in an economic benefit of offloading Internet traffic from cellular onto V2X networks. Our previous work has shown that deploying V2X infrastructure for offload is cost-effective in urban areas [5]. This will likely be relevant for the foreseeable future. Although macrocellular capacity continues to increase as carriers expand infrastructure and regulators allocate more spectrum, it is unclear whether cellular capacity will increase dramatically in the near future for highly mobile users. Moreover, mobile Internet traffic has grown 18-fold in the past 5 years [6], justifying alternative approaches such as data offload. We have also shown that it is even more cost-effective if infrastructure is shared between Internet access and safety applications [7]. However, the work in [5], [7]–[9] considered the bandwidth allocated for ITS as fixed and not shared. In contrast, this paper focuses on spectrum management; we examine the economic benefit of adding spectrum to offload Internet traffic. If the marginal benefit of adding one unit of spectrum exceeds its opportunity cost (i.e., the foregone benefit of using that spectrum for something else), then that unit is worth allocating for ITS. With this approach, we estimate the ITS bandwidth that maximizes benefit minus cost. In addition, we examine how that estimate changes with uncertain factors such as data rates of Internet traffic and penetration of V2X devices in vehicles. On the issue of whether ITS spectrum should be shared with unlicensed devices, one question is what the difference between throughputs to vehicles and unlicensed devices in spectrum exclusively allocated, and throughputs in shared spectrum. Another question is how much spectrum is needed to carry a given amount of data from vehicles and unlicensed devices when each type of device uses separate spectrum, and how much spectrum is needed to carry the same data if the spectrum is shared. A third question we address is what sharing scheme should be adopted, if any.

The debate over the FCC NPRM is primarily about which spectrum-sharing scheme causes less interference to safety-related communications. In contrast, we consider a scenario where part of the ITS band is allocated for safety messages and not shared with other types of communications, but the rest of the spectrum is shared between V2X and unlicensed devices for non-safety communications, on a co-equal basis. In that scenario, Internet Service Providers (ISPs) deploy V2X-based roadside units (RSUs) connected to the Internet to offload part of the vehicular traffic. In addition, ITS spectrum allocated for non-safety ITS traffic can be shared with unlicensed devices such as Wi-Fi hotspots.

II. RELATED WORK

To the best of our knowledge, this is the first work that examines whether it is cost-effective to allocate more or

Manuscript received November 13, 2018; revised March 28, 2019; accepted April 29, 2019. Date of publication May 10, 2019; date of current version September 9, 2019. This project is funded in part by Carnegie Mellon University's Mobility21 National University Transportation Center, which is sponsored by the US Department of Transportation, in part by the National Science Foundation under Grant No. 1547237, and in part by the CMU-Portugal Partnership and Portugal Foundation for the Science and Technology. The associate editor coordinating the review of this paper and approving it for publication was S.-L. Kim. *(Corresponding author: Alexandre K. Ligo.)*

The authors are with the Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: aligo@cmu.edu; peha@cmu.edu).

Digital Object Identifier 10.1109/TCCN.2019.2916026

less ITS spectrum than what is currently allocated, based on economic marginal benefit and marginal cost of that spectrum.

There is previous research on issues related to sharing of ITS spectrum with unlicensed devices. Such sharing can be done in several ways [10] that include the leading schemes proposed [1], [11] in response to the FCC's sharing NPRM, and other possible arrangements. Previous work mostly examined harmful interference to safety messages on those proposals [1], [12]–[14]. While these works argued that safety applications can be significantly impacted by sharing, some offered regulatory measures for device and protocol improvements to mitigate such interference [12], [13].

Whereas interference on V2X safety has been subject of extensive research and testing, we assumed that safety communications will use dedicated spectrum, and focused our work on the portion that can be shared for non-safety traffic. Some papers investigate issues similar to the ones we do. Reference [15] studied the performance of Wi-Fi devices in the ITS band, using a testbed of two DSRC and two Wi-Fi nodes. They found that both V2X and Wi-Fi performance can degrade with sharing for certain conditions, but changes in protocol parameters and channelization can mitigate such degradation. The work in [16] compared throughput capacity to V2X and Wi-Fi devices among different sharing scenarios, including scenarios of separate spectrum and sharing among equals. They conclude that sharing can result in significant improvement in throughput capacity for unlicensed devices, while causing "acceptable" degradation in V2X performance (in the order of 10% or less). This is related to our work because our investigation of whether to share includes an analysis of throughput performance with and without sharing. However, the research questions, scenarios and contributions are different. First, their assumptions and results are more applicable to DSRC for safety communications, while ours apply to Internet traffic. This is because they consider DSRC nodes broadcasting data, which is typical for safety applications. In contrast, we consider unicast connections over a mesh network among vehicles and Internet-connected RSUs, over which TCP connections carry Internet traffic. Second, they assume that the locations of DSRC and unlicensed devices are placed according to Poisson point processes (and show that results differ from those in realistic locations), while we derive locations of devices from vehicles, residential hotspots, and road locations from a real city. Another difference is that they use theoretical channel capacity to compare different sharing schemes, each with a fixed amount of spectrum. We instead determine data throughput resulting from existing protocol mechanisms (including, e.g., collisions, TCP flow and congestion control), for varying amounts of spectrum, to find the amount of spectrum used in different sharing schemes. Moreover, while their study is relevant to some of the questions addressed in this paper, we also address novel questions such as how much ITS spectrum to allocate, based on its benefit and cost. Therefore, our work significantly differs from existing research on sharing of the ITS band.

III. Method

We employ an engineering-economic approach, of which a major part is to use packet-level simulation to examine how the ability of vehicles and hotspots to carry IP traffic is affected by sharing spectrum between those types of devices. To define simulation parameters, data from a real vehicular network and Wi-Fi hotspots operating in Portugal is used. Several factors are varied. One is the amount of spectrum, for which we assume either that vehicles and hotspots use separate spectrum, or spectrum is shared with vehicles and hotspots coexisting in a co-equal basis, using 802.11 listen-before-talk mechanisms. Other factors that are varied include the densities of vehicles and hotspots, data rates of incoming Internet traffic, and whether hotspots are indoors or outdoors. Another part of our method addresses the issue of how much spectrum to allocate for ITS, by using the vehicular simulated throughput to estimate the economic benefit of adding ITS spectrum to offload Internet traffic.

A. Model of Usage and Sharing of the ITS Band

The answers to how much spectrum to allocate for ITS and whether it should be shared depend on benefits accrued by using the ITS band by V2X and unlicensed devices. In this subsection we describe the assumptions regarding the use of the ITS band and what benefits are considered.

Among the leading schemes proposed for the FCC NPRM on spectrum sharing [4], one is based on unlicensed devices being allowed to use only part of the ITS band, while the other part is reserved for safety traffic and not shared [11]. In the shared channels, DSRC devices and unlicensed devices would coexist on a co-equal basis, which means that the proposal is not to grant priority access, but rather allow DSRC and unlicensed devices to coexist in shared spectrum through mechanisms such as "listen before talk."

Like the coexistent sharing-among-equals proposal, in our model safety messages are transmitted exclusively over dedicated channels where no other type of traffic is allowed. We assume those dedicated channels are sufficient to carry all safety traffic, and no additional safety benefit is achieved if spectrum is allocated beyond the dedicated channels [11]. This model allows us to evaluate non-safety benefits from adding spectrum, in a way that is independent from safety benefits. In our model, using spectrum not dedicated for safety produces the benefit of carrying Internet traffic either to V2X, to unlicensed devices, or both. We assume spectrum is used for IP traffic as follows. Bidirectional connections are established between each vehicle equipped with a V2X onboard unit (OBU) and one RSU which serves as an Internet gateway. A vehicle can connect to an RSU either directly or through multiple hops with other vehicles acting as relays (Fig. 1). Each vehicle uses one channel from $D \geq 0$ channels, while each RSU can use all D channels. Each hotspot uses one channel from either S channels ($0 \leq S \leq D$) that are shared with V2X devices, or $W > 0$ channels located in a separate band.

In our model, devices share channels through the listen-before-talk mechanisms specified in IEEE 802.11 as co-equals,

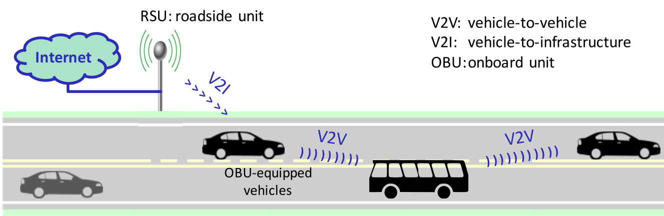

Fig. 1. V2X-based communications with V2V and V2I links.

i.e., vehicles and hotspots have equal priority when transmitting. V2X and unlicensed devices detect each other's transmissions by preamble detection, which means devices are able to decode the packets of the other type of device.

In addition to those common rules, spectrum sharing takes place according to one of three possible sharing schemes, which differ on whether devices of different types cooperate with each other [10] and how. We call the simplest scheme *coexistence* without cooperation, where V2X and unlicensed devices sense each other's transmissions, but devices of one type avoid interference without explicitly cooperating with devices of the other type. In the two other schemes hotspots cooperate to relay V2X packets, in order to be allowed in shared spectrum [17]. In one of the cooperation schemes, vehicles out of communications range with RSUs or other vehicles can send packets to hotspots, which relay those packets to and from their backhaul connections to the Internet. We call this scheme *backhaul cooperation*. In the second cooperation scheme, hotspots relay traffic between V2X devices, i.e., a vehicle that is not in range with RSUs or other vehicles can use hotspots as non-moving relays in a multihop route to reach an RSU. We refer to this as *relay cooperation*.

We assume 10 MHz channels, which is the current specification in the ITS band. One channel is used for each hotspot connection, and for each hop in a vehicle-RSU connection. Also, devices choose channels to transmit before establishing the connections. Channels used by vehicles are chosen according to [18], which considers the expected interference from nodes already assigned to the channels. For hotspots, one channel is assigned at random per hotspot, with all channels having equal assignment probability.

B. Economic Benefit of ITS Spectrum for Internet Access

Our model assumes that a dedicated portion of ITS spectrum provides all safety benefits, and the factor that determines how much spectrum to allocate for ITS is the marginal benefit per MHz of carrying Internet traffic over V2X. We assume that in the absence of V2X, in-vehicle Internet is served by macrocellular networks. Such traffic is expected to increase rapidly and exponentially over time [6], so infrastructure capacity must be expanded at this rate where the cellular network is capacity-limited. However, when Internet traffic is carried over V2X at peak hours fewer macrocellular towers are needed than in a scenario without V2X. If the "avoided" cost from fewer towers exceeds the cost of V2X, then deploying V2X is cost-effective and enhances social welfare when compared to expanding cellular infrastructure. Hence, we define the benefit of offload as the cost savings from deploying fewer cell towers.

The total number of towers "saved" per km^2 is given by

$$\rho_{savedtowers} = \frac{bpsOff \;\; FR}{s_{sector}\, bw N_{sectores}} \tag{1}$$

where *bpsOff* is the peak-hour, downstream V2X throughput per km^2, *FR* is the frequency reuse factor, s_{sector} is the average downstream spectral efficiency in bps/Hz/sector, *bw* is the total downstream bandwidth per sector per cellular carrier, and $N_{sectors}$ is the number of sectors per tower.

Social welfare is maximized when the marginal benefit per unit of spectrum added equals the marginal costs of offloading [19]. Costs can be of three types. One is the opportunity cost of not allocating the spectrum for a use other than ITS. The opportunity cost is the economic surplus that would be obtained in the best use of the spectrum other than ITS.

The second cost is of RSUs deployed for Internet access. The density of RSUs deployed for Internet access N_{RSU} affects total RSU cost and also benefit through *bpsOff*. For fixed bandwidth and other factors, we choose N_{RSU} as the RSU density that maximizes benefit minus RSU cost. N_{RSU} is determined for each scenario of numerical assumptions. Moreover, our simulations suggest that in a given scenario N_{RSU} is approximately insensitive to bandwidth. Hence, we keep N_{RSU} fixed when spectrum amount is varied in a scenario. The modeling of RSUs cost is described in [5], [7], [8].

In addition to marginal benefits and costs we also examine average benefits and costs in each scenario. This is because it is possible that for a certain bandwidth marginal benefit equals or exceeds marginal cost, but average benefit does not, since RSU cost is an upfront cost that can be higher than benefit.

The third cost is of OBUs. As a base assumption, we consider the case where there is a mandate to equip cars with OBUs for safety, as may occur in the U.S. [20]. In this scenario OBU costs are incurred for safety and thus do *not* matter for non-safety purposes. (In the results Section we discuss possible implications when a mandate does not occur.)

C. Opportunity Cost of Spectrum Allocated for ITS

The cost of spectrum at 5.9 GHz is uncertain, but we can use available evidence to estimate an upper bound. In the case of spectrum allocated for licensed use, a popular way of estimating its opportunity cost is to use the prices paid in license auctions. In recent U.S. auctions, winning bids exceeded \$2 per unit of spectrum per capita (MHz-pop) for bands in 1.8-2.2 GHz in 2015 [21]. It must be taken into consideration that spectrum is far less valuable at higher frequencies [22], [23], perhaps by an order of magnitude. That might place the value of ITS spectrum in the order of a few tens of cents. However, newer technology operates effectively at higher frequencies than was typical in the past, so spectrum value is probably changing, which adds to its uncertainty.

Moreover, the FCC NPRM on sharing indicates that ITS spectrum might be opened for unlicensed use. Estimating the marginal value of unlicensed spectrum is difficult, but marginal value per MHz would certainly be less than value per MHz averaged over all spectrum. A group interested in expanding unlicensed use of spectrum has estimated the total value of [24] would average about \$0.70 per MHz-pop. Therefore, the opportunity cost is likely well below this value, perhaps in the vicinity of \$0.20-\$0.40 per MHz-pop.

D. Simulation Model and Assumptions

Our method depends on estimates of throughputs, which we simulate at packet-level using ns-3. The part of the simulation model that represents vehicles and RSUs is described in detail in [5], [7]–[9]. We have extended the model to vary the amount of ITS spectrum and to allow sharing with unlicensed devices. The network is simulated with vehicles changing positions each 5 seconds. During a 5-s interval, throughput is simulated over a network of non-moving nodes. Then vehicle positions are changed and the process is repeated. Vehicles are positioned every 5 s according to the GPS logs of buses and taxis in Porto. Each vehicle connects to an RSU through TCP/IP. Packets are routed through connections with up to three hops. The received signal is calculated according to [25] (urban microcell B1), in which distance affects signal-to-noise-and-interference ratio, and therefore link speeds between vehicles and RSUs.

The assumptions for the Wi-Fi traffic are as follows. We adopt the simplifying assumption that all traffic to a hotspot is carried through a single TCP connection between the hotspot and a client device located 10 m away. We consider both indoor and outdoors hotspots. For indoor hotspots, some are active while others are not. Active hotspots are receiving packets at a constant rate throughout the 5 s. Every 5 s, a different set of hotspots is randomly selected to be active. Moreover, the density of indoor hotspots in an area depends on population density, and their positions for the simulation are randomly sampled from the set of coordinates obtained from the Wi-Fi provider in Porto (see Section III-E below). If the quantity of coordinates to be used in a simulation is higher than the total number of coordinates in the dataset, then the coordinates that exceed the total are also sampled from the same set as follows. One neighbor hotspot is randomly selected from the three closest neighbors of the hotspot to be shifted. Then its new position is chosen randomly between the original position of the hotspot and the position of its neighbor. All hotspots have a height of 3 m. This overstates the interference where hotspots are far from the ground in multi-story buildings. The signal transmitted by a hotspot propagates according to an indoor propagation model [25] to the endpoint of its TCP connection, or a model with wall obstruction to V2X devices or outdoor hotspots.

The assumptions for outdoor hotspots are different. These are placed along the streets of Porto (see Section III-E below) with a fixed inter-hotspot distance. Link speeds between hotspots and vehicles also depend on distance, since these links are subject to the same loss model used for vehicles. Moreover, all outdoor hotspots are active at peak hours. The transmission power of 11 dBm at the antenna output is consistent with popular Cisco Wi-Fi hotspots [26].

The number of channels D and S (or W), and the selection of the channel used by each node is defined before the simulation of a 5-s interval is run.

E. Portugal Dataset

To set some of the simulation parameters, we use data from a real vehicular network, Wi-Fi hotspots and the coordinates of roads in Porto, Portugal [5], which we believe are representative of urban areas in any industrialized country. We used a dataset with measurements of data transferred over V2X and cellular, and GPS position data of 800+ buses and taxis. Also, we have collected positions of 65,000+ Wi-Fi hotspots in Porto, which were available in the website of FON, one major Wi-Fi provider who partners with a major fixed broadband provider in Portugal. Therefore, the data is probably representative of hotspots in households and small businesses. We also use the coordinates of city roads. Porto data is used in four ways. First, GPS positions are used to determine the positions of the vehicles in the simulation as described in Section III-D. Second, strength of the signal received from RSUs is measured in the buses, and is verified to be compatible with the simulated signal, on average. Third, coordinates of the Wi-Fi hotspots are used to determine the positions of indoor hotspots. Fourth, road locations are used to determine the positions of outdoor hotspots in the simulation as described in Section III-D.

F. Base Case Numerical Assumptions

Table I shows the base case numerical assumptions used for the calculation of benefits and costs as in Section III-B. These assumptions apply for the results in the following section unless otherwise stated. (For justification of the numerical assumptions such as the number of V2X-equipped vehicles per population density and other values, see [5] and [7].)

Many of the results presented refer to a penetration of V2X in vehicles of 100%. This is reasonable over the time-frame of a spectrum allocation decision if the Department of Transportation mandates V2X for safety [27]. Another assumption that is highly uncertain is the data rate per vehicle. We assume a "low" case value of 400 kbps [5] and a "high" case value of 4 Mbps, because data rates have been increasing rapidly [28]. (e.g., typical video streams have an average bitrate of 2 Mbps [29], with Netflix HD encoding.)

For other values, we use base assumptions that are representative of five years into the future (see, e.g., [28]). Although this work informs spectrum allocation decisions that may span decades, the rate of technological change makes decade-long predictions highly uncertain.

We considered two different values for average data rate in the peak hour of active hotspots in five years. Our low estimate is 5 Mbps. The majority of traffic in the U.S. is from video applications [30], and this 5 Mbps value would be appropriate if each hotspot supported one video stream throughout the peak hour, and if data rates increased from the current average at 19% per year for 5 years [28]. Our high estimate is 27 Mbps (see Table I).

For indoor hotspots, we assume that 15% of them will be active at a time. Current estimates for the average traffic in U.S. households are currently around 100 GB per month [28], [30]. A hotspot transferring 300 kbps at all times would transfer 100 GB over a month, then the share of active

TABLE I
BASE CASE ASSUMPTIONS (SEE [5], [7] FOR DETAILS)

Assumption	Value
Discount rate	7%, real
Time horizon for analysis	5 years
Data rate of incoming Internet traffic per vehicle on the road	Low scenario: 400 kbps average - 50% of cars are endpoints for 800 kbps, 50% are relays only. High scenario: 4 Mbps average
Share of downstream traffic	90% of data from RSU to vehicle
Macrocellular spectrum efficiency s_{sector}	1.4 bps/Hz/sector (downstream average)
Sectors per macrocell $N_{sectors}$	3
Macrocellular bandwidth bw	70 MHz (downlink per cellular carrier)
Reuse factor FR	1 (macrocellular frequency reuse)
Unit cost of macrocellular tower C_{tower}*	$750,000 (see e.g. [31]): NPV of capital and operating expenses (Capex and Opex) over time horizon
Cost of one V2X Internet-only RSU c_{io}*	$14,000 (NPV of Capex and Opex over time horizon)
Population density	5,000 people/km^2
Indoor hotspots per capita	1 hotspot for every 4 people (see section D) [6]
Outdoor hotpot locations	Placed along roads every 150 m [32]
Data rate of incoming Internet traffic per hotspot	Low scenario: 5 Mbps at 15% of hotspots at each 5-s interval (see below). High scenario: 27 Mbps

* Monetary values in 2014 U.S. dollars

hotspots is assumed as 300 kbps / 2 Mbps = 15%. Although this assumption about the share of active hotspots at a given time into the future is uncertain, it is likely that not all active hotspots would be using the channels in the ITS band in any given time. Hence, this assumption may result in conservative throughput estimates, given that the real interference from indoor hotspots may be lower than what we estimate.

For outdoor hotspots, we assume they to be placed every 150 m in all urban roads. Since deployment of outdoor Wi-Fi has been limited to a few downtown areas and other sparse locations, this assumption is also likely to result in higher interference to vehicles than in typical urban areas. (For this reason, we compare scenarios with both indoor and outdoor hotspots with scenarios with indoor hotspots only.)

For the results of a specific location, the base-case population density is 2,000 people/km^2, which represents a city like Pittsburgh, unless stated otherwise.

IV. RESULTS

We address the issue of how much spectrum to allocate for ITS, whether the ITS band should be shared with Wi-Fi devices, and if so, what sharing scheme should be used. The throughput for each scenario of bandwidth, device density and data rates is derived by averaging throughput for at least 1000 vehicles. Assuming that the throughputs of the vehicles are mutually independent, then the 95% confidence interval is within 5% of the mean throughput.

A. How Much Spectrum to Allocate for ITS

In this subsection we estimate economic benefits and costs of deploying V2X infrastructure for Internet access on a nationwide scale for the U.S. For this estimate, we assume spectrum is used for ITS only, i.e., it is not shared with unlicensed devices. We then use benefits and infrastructure costs to derive the bandwidth that maximizes social welfare as a function of uncertain factors such as the opportunity cost of spectrum in the ITS band, data rates, and OBU penetration.

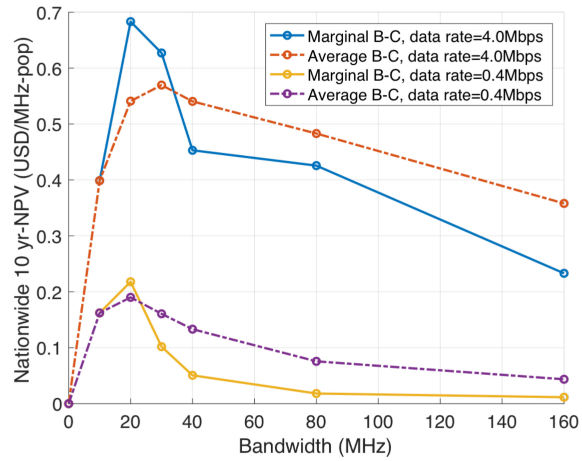

Fig. 2. Nationwide benefit minus RSU cost per capita (B-C), as a function of bandwidth. Lines for two different data rates of incoming traffic per vehicle are shown. OBU penetration is 100%.

We quantify economic benefits and costs of allocating a given amount of spectrum for ITS throughout the entire nation, even in regions where population density does not justify V2X networks (i.e., where there is no benefit but there is a cost of spectrum), because this is generally how spectrum is allocated. We calculate benefits and costs of using the spectrum for ITS in each U.S. census tract and then sum benefits and costs over all tracts. We assume that RSU deployment decisions are made at the census tract level, i.e., N_{RSU} (see Section III-B) is determined at each census tract based on its average population density (this was employed in [31], [33]).

Fig. 2 shows marginal and average benefit minus RSU cost (B-C) per MHz-pop (nationwide - U.S.), as a function of bandwidth allocated exclusively for vehicles. It shows results for two data rates of incoming Internet traffic per vehicle (low and high scenario as in Table I). The other assumptions are base case values. For a particular bandwidth to be worth allocating, both marginal and average B-C must exceed the opportunity cost of ITS spectrum [19]. If marginal B-C is less than cost at bandwidths where marginal B-C decreases with bandwidth, then reducing bandwidth results in higher B-C. If average B-C is less than cost, then B-C is greater with a bandwidth of 0. Fig. 2 shows that B-C does not change monotonically with bandwidth. This is because while marginal and average benefit do decrease monotonically with bandwidth, RSU cost does not. It is proportional to the number of RSUs. However, the quantity of RSUs deployed is roughly invariant with bandwidth higher than 20 MHz, for the range of data rates, densities and other factors we considered.

Given the uncertainty in the opportunity cost of ITS spectrum (Section III-C), we examine the relationship between the opportunity cost and the optimal bandwidth. For a given opportunity cost, Fig. 3 shows the maximum bandwidth for which marginal and average B-C exceed that cost. In Section III-C we conjecture that the cost of ITS spectrum might be around

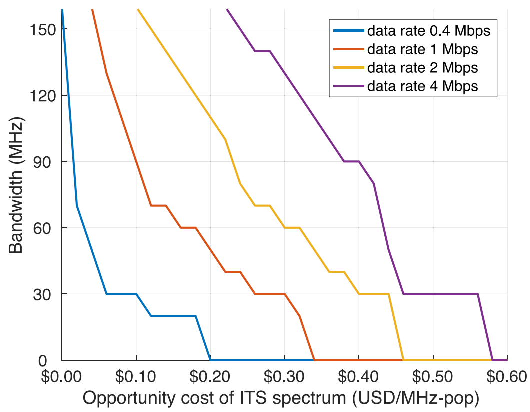

Fig. 3. Bandwidth that maximizes social welfare (nationwide), as a function of the opportunity cost of ITS spectrum. Curves are shown for distinct data rates of incoming traffic per vehicle. OBU penetration is 100% and other numerical assumptions are in base case values.

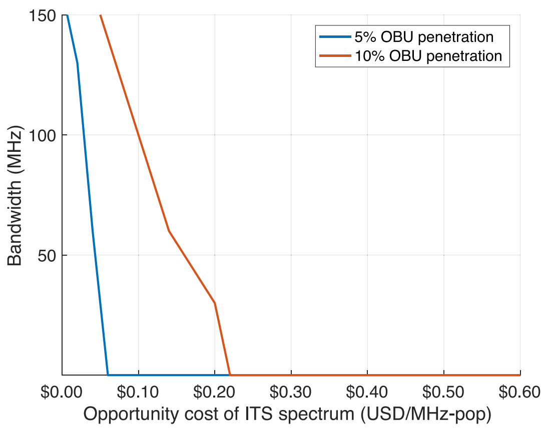

Fig. 4. Bandwidth that maximizes social welfare (nationwide) as a function of the opportunity cost of ITS spectrum. Curves are shown for distinct penetrations of OBUs. Data rate per OBU is 27 Mbps (maximum for 802.11p in 10 MHz channel) and other numbers are at base case values.

$0.20-$0.40 per MHz-pop. Fig. 3 shows that for such costs it might be worth allocating spectrum, but the amount that maximizes social welfare depends not only on cost but also on other factors as well. For example, at an OBU penetration of 100% of vehicles and average data rate of incoming traffic of 4 Mbps/vehicle, Fig. 3 shows that it is worth allocating 40 MHz of ITS spectrum, which is the bandwidth currently available for non-safety, as long as the opportunity cost of spectrum is below $0.45 per MHz-pop. However, for a lower average data rate of 0.4 Mbps per vehicle the same bandwidth could be allocated only if cost is much lower (below $0.05 per MHz-pop). From Fig. 3 we see that there are realistic scenarios in which it is worth allocating more spectrum than it is currently available for non-safety use, but there are also scenarios in which it is worth allocating less, if any.

Moreover, the results above are for an OBU penetration of 100%, which is consistent with a mandate of V2X. Out of the context of a mandate, lower penetrations are possible, with OBUs more likely in vehicles that demand higher data rates. Fig. 4 shows the bandwidth that maximizes social welfare in such a scenario and shows that bandwidth is highly sensitive to penetration. The range of opportunity costs that results in any bandwidth to be allocated is significantly smaller in Fig. 4 than for the scenarios with 100% penetration (Fig. 3). However, Fig. 4 shows that a small increase in penetration (5% to 10%) changes significantly the bandwidths worth allocating. For example, at 10% penetration it is worth allocating 40 MHz (the bandwidth currently available for non-safety) if the cost of spectrum is $0.18/MHz-pop. However, a scenario where it is not worth allocating spectrum in excess of safety is also plausible for low OBU penetrations or if spectrum costs more than a few tens of cents per MHz-pop.

It is important to note that this discussion applies for spectrum allocated exclusively for ITS, which does not capture the value of sharing spectrum with unlicensed devices. Benefits of sharing are discussed in the following subsection.

B. *Should ITS Spectrum Be Shared With Unlicensed Devices?*

To examine whether and how to share spectrum, we compare performance of four different strategies: allowing vehicles and unlicensed devices to share spectrum with the three possible sharing schemes defined in Section III-A, and placing vehicles and unlicensed devices in separate bands. That comparison is done with respect to two measurements. The first is the estimated throughputs to vehicles and unlicensed devices when different types of device uses separate channels, and examine how those throughputs differ when the devices use shared spectrum. The second measurement is the amount of shared spectrum needed to achieve given throughputs, which is compared to the total amount of spectrum in separate bands to achieve the same throughput.

Fig. 5 and 6 show vehicle throughput and hotspot throughput. In both graphs the horizontal axis is the bandwidth allocated (in excess to what is used by safety). We show throughputs for indoor hotspots only and indoor+outdoor hotspots. As expected, throughput increases with bandwidth for both vehicles and hotspots, although at a diminishing rate.

Also, throughputs differ among strategies. Fig. 5 shows that the difference between throughput to vehicles on exclusive spectrum and throughput on shared spectrum is negligible in some scenarios and significant in others, depending on factors such as whether outdoor hotspots are present or not. For example, differences between the curves in Fig. 5 suggest that vehicle throughput is significantly affected by outdoor hotspots. With indoor hotspots only (left graph), the differences between throughputs on exclusive spectrum and throughput on shared spectrum is within the 95% confidence interval, regardless of sharing strategy and bandwidth. In this scenario, most vehicles are not close enough to a hotspot to experience or cause harmful interference at any given time. Besides, indoor hotspots are separated from streets by walls and thus cause low impact on vehicle throughput. However, the difference in throughputs among strategies can be high when outdoor hotspots are present (Fig. 5, right). Throughput to vehicles coexisting with unlicensed devices is significantly lower than throughput without sharing. The loss of vehicle throughput

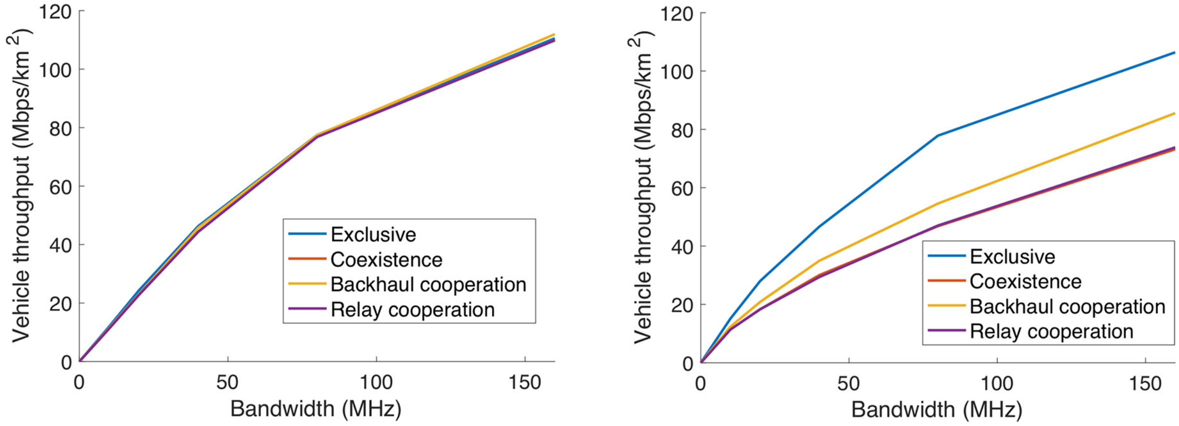

Fig. 5. Vehicle throughput as a function of spectrum allocated, under different strategies. The left graph refers to indoor hotspots only, the right graph is for indoor and outdoor hotpots. The other assumptions are at base-case values (with "high" data rates).

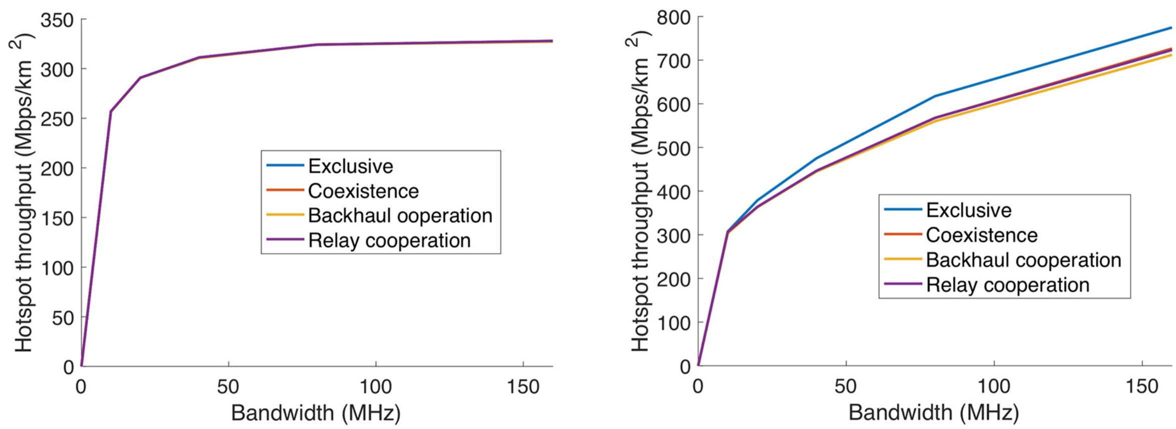

Fig. 6. Hotspot throughput as a function of spectrum allocated, under different strategies. The left graph refers to indoor hotspots only, the right graph is for indoor and outdoor hotpots. The other assumptions are at base-case values (with "high" data rates).

caused by sharing can be mitigated if unlicensed devices are required to cooperate, as shown in Fig. 5 (right). One reason is that there are vehicles that are not in communications range of an RSU. (It would not be cost-effective to deploy RSUs ubiquitously.) Hence, if hotspots help some of those disconnected vehicles reach the Internet, overall vehicle throughput is higher. However, not all cooperation schemes increase vehicle throughput (relative to coexistence). Fig. 5 shows that backhaul cooperation results in higher vehicle throughput when there are outdoor hotspots (right). However, relay cooperation does not result in change.

Fig. 6 shows that the difference between hotspot throughput in separate spectrum and in shared spectrum is small for all sharing schemes. For all scenarios simulated the difference is less than 10%, which is mostly within the 95% confidence interval, and is well below that value for most scenarios. The impact of sharing on hotspot throughput is small because most vehicles are not close enough to a hotspot to experience or cause harmful interference at any given time. The densities and data rates of hotspots may be much higher than that of vehicles, but many of them are indoors.

We have found that the presence of outdoor hotspots has significant impact on vehicle throughput. Hence, in Fig. 7 we examine throughputs for varying densities of outdoor hotspots. The horizontal axis shows decreasing distance between outdoor hotspots (the base-case value is 150 m). The left graph shows that vehicle throughput is higher with backhaul cooperation than with coexistence. The relative difference increases as hotspot density increases from 0 to 1/300 m^{-1}, as more vehicles that were unable to reach an RSU can now reach a hotspot. This difference in throughput remains constant for greater densities, because nearly all vehicles have been connected. On the other hand, vehicle throughput with relay cooperation is not significantly different from the throughput with coexistence, for all densities of outdoor hotspots examined.

Fig. 7 (right) shows that hotspot throughput with all sharing strategies is less than throughput with exclusive spectrum, for all hotspot densities. Moreover, the curves for the sharing schemes all overlap, indicating that there is little difference in the burden that these strategies impose on hotspot throughput.

We also analyze how throughputs vary with other factors, such as data rates and densities of devices other than outdoor hotspots. For all those scenarios, the differences between the throughputs on shared spectrum and the throughputs on separate spectrum are not as large as the scenarios of Fig. 7. Moreover, the differences in throughputs among the different sharing strategies are similar to those found in Fig. 7.

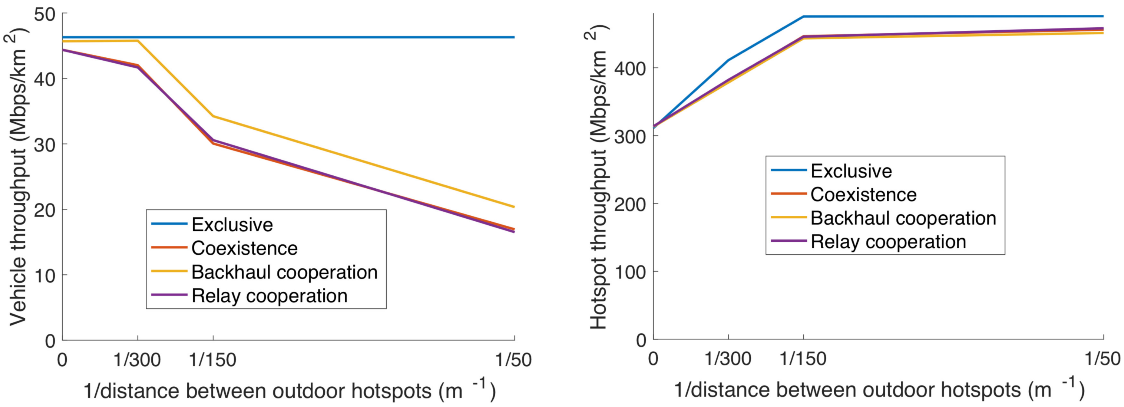

Fig. 7. Throughputs to vehicles (left) and hotspots (right) as a function of decreasing distance between outdoor hotspots, under different strategies. Bandwidth is 40 MHz, and the other assumptions are at base-case values (with "high" data rates).

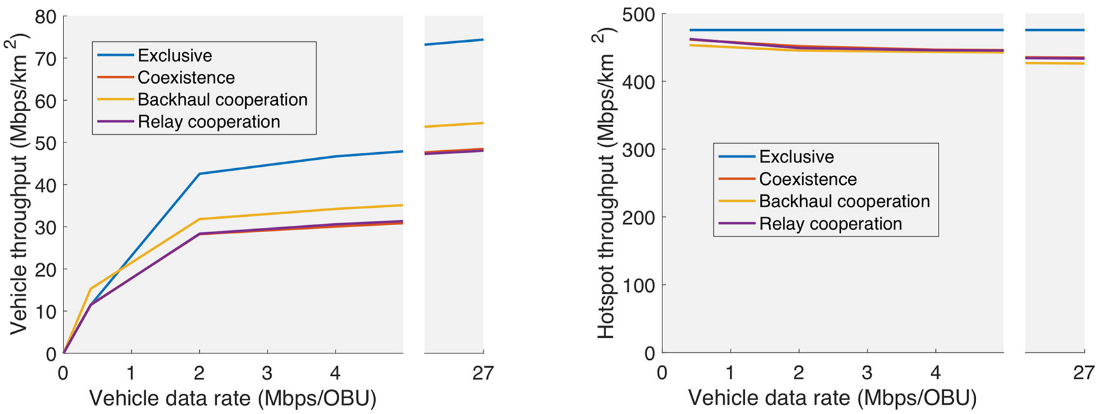

Fig. 8. Throughputs to vehicles (left) and hotspots (right) as a function of data rates of Internet traffic per vehicle, under different strategies. Bandwidth is 40 MHz, and the other assumptions are at base-case values.

Fig. 8 shows throughputs for varying data rates of vehicular Internet traffic. The left graph shows that for the "low" data rate (0.4 Mbps/vehicle), vehicle throughput with coexistence is the same as without sharing. This is because at this data rate vehicles face less congestion in the channels than at higher data rates. Nevertheless, backhaul cooperation achieves higher throughput than relay cooperation and coexistence, because there are vehicles communicating via backhaul cooperation that were disconnected from RSUs in the other strategies.

As data rate increases, Fig. 8 (left) shows that vehicle throughput with any sharing strategy is significantly less than throughput without sharing. This is because at higher data rates, vehicles face more congestion when sharing channels than on exclusive spectrum. Still, backhaul cooperation results in higher throughput than coexistence, while relay cooperation results in roughly the same throughput as coexistence.

Fig. 8 (right) shows that throughput to unlicensed devices with higher vehicle data rates is less than throughput with lower data rates, because of the higher interference from vehicles and RSUs. However, that difference is relatively small, because of the smaller quantity of vehicles compared to the quantity of hotspots. As with previous graphs, Fig. 8 (right) also shows that hotspot throughput with all sharing schemes is slightly less than throughput with exclusive spectrum, and the hotspot throughput at any scheme is not significantly different from the other sharing schemes.

Like data rates, the penetration of V2X devices in vehicles is also expected to increase over time, resulting in higher densities of V2X devices. The effect of increasing V2X penetration on throughputs for different sharing strategies is similar to those previously shown for increasing data rates.

Backhaul cooperation results in higher vehicle throughput than coexistence, for a wide span of conditions of data rates and device densities. However, the difference between sharing schemes may disappear for extremely high densities of devices. Fig. 9 (left) shows the effect of increasing population density (which results in more of both vehicles and hotspots) on vehicle throughput. Although backhaul cooperation results in higher vehicle throughput than with coexistence for most population densities, the difference diminishes for more than 7,000 people/km^2, because the quantity of unlicensed devices increases faster with population density than the quantity of vehicles increases, so the ratio of hotspots to vehicles is higher with higher population density. Hence, vehicles face more interference from hotspots. However, this density is extremely high. (Few locations have 7,000 people/km^2 or more.) The differences between the sharing strategies shown for densities below 5,000 people/km^2 should hold for most

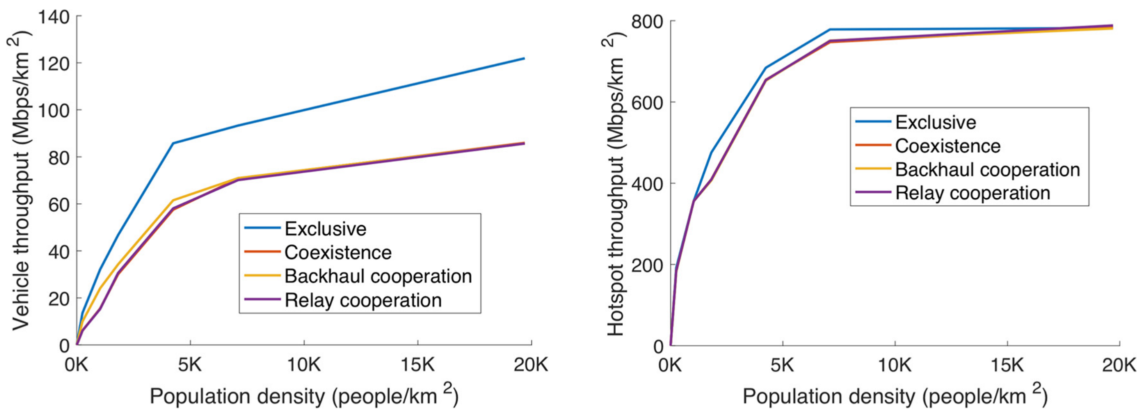

Fig. 9. Throughputs to vehicles (left) and hotspots (right) as a function of population density, under different strategies. Bandwidth is 40 MHz, and the other assumptions are at base-case values (with "high" data rates).

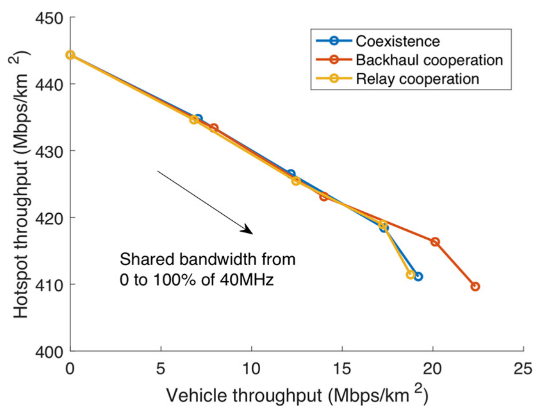

Fig. 10. Hotspot throughput as a function of vehicle throughput, with different sharing schemes. Bandwidth is fixed at 40 MHz, and the other assumptions are at base-case values (with "high" data rates).

of the U.S. Fig. 9 (right) shows that for 1,000 people/km^2 or less, there is no significant difference between hotspot throughput with sharing and throughput on separate spectrum. This is because of the smaller quantities of devices sharing the channel. Most of the U.S. has population densities below 1,000 people/km^2. With higher population density, hotspot throughput with sharing is slightly less than throughput with exclusive spectrum (although the difference is statistically significant). However, this difference is negligible for 7,000 people/km^2 or more. Hotspot throughput is not significantly different among sharing strategies.

For a sharing scheme to be the best, it must achieve throughputs for the two device types that other schemes cannot. Fig. 10 shows the throughputs that can be achieved when part of the spectrum is shared using a given sharing scheme, and the rest of the spectrum is available only to one of the two device types. The total amount of spectrum is kept fixed, and the amount of spectrum shared is varied. Thus, given throughputs to vehicles and unlicensed devices can be achieved if and only if the point associated with those two throughputs falls within the feasible region [34], which is the region bounded by the curve associated with that sharing scheme and the X and Y axes. The larger the feasible region, the better. The graph shows that the edges of the feasible region for all sharing strategies overlap if vehicle throughput is less than about 17 Mbps/km^2. For that range of the graph, the same vehicle and hotspot throughputs can be achieved with any of the three sharing schemes. However, the feasible region of backhaul cooperation is larger than the regions of the other schemes. There is a range of vehicle throughput (between 17 and 19 Mbps/km^2, for the assumptions used) that can only be achieved with backhaul cooperation, or by increasing the total amount of spectrum used.

The results shown so far are related to previous research because we investigate whether to share with an analysis of throughput with and without sharing. However, our research questions, scenarios addressed, and results presented in this section are different. First, the assumptions and results from work such as [12]–[16], [35]–[37] are more applicable to DSRC for safety, while our work applies to Internet traffic. This is because they consider data broadcasting data, which is typical for safety applications, while we consider unicast connections over a mesh of vehicles and Internet-connected RSUs, over which TCP/IP connections. Second, previous work such as [16] assume that the locations of DSRC and unlicensed devices are placed according to random distributions such as Poisson point processes (and the authors show that results differ from those derived from realistic locations), or a handful of devices places arbitrarily. Rather, we derive locations of V2X and unlicensed devices from vehicles, residential hotspots, and road locations from a real city. Another difference is that [16] use theoretical channel capacity to compare different sharing schemes, each with a fixed amount of spectrum. We instead determine data throughput resulting from protocol mechanisms including, e.g., collisions, TCP flow and congestion control, for varying amounts of spectrum.

Despite the differences, most of previous work found that Wi-Fi devices can interfere significantly with the performance of DSRC devices, especially under high load when, e.g., every Wi-Fi device always have a packet to send. However, most authors also concluded that if the interval between Wi-Fi packets is above a certain threshold, performance of DSRC approach that in the absence of sharing. Although the scenarios and assumptions are very different from ours, their conclusions are consistent with our findings regarding how each type of device affects the other performance.

We also look into how much spectrum is needed to carry a given amount of data from vehicles and unlicensed devices over separate channels, and how much spectrum is needed to carry the same amount of data on shared spectrum, for different sharing schemes. We determine those amounts of shared and separated spectrum as follows. First, we find vehicle hotspot throughputs for a given amount of shared spectrum and a given sharing scheme. Then, we find the amount of spectrum used to achieve that same vehicular throughput, but on spectrum used by vehicles only. Likewise, we find the amount of spectrum used by hotspots only. The process is repeated for several throughputs.

Fig. 11 shows the amounts of spectrum obtained with the procedure above as a function of vehicle and hotspot throughputs. One curve is the total amount of spectrum when vehicles and hotspots use spectrum separately, and the others show the amount of spectrum with different sharing schemes. The curves for any given vehicle throughput also refer to the same hotspot throughput (i.e., the curves at any given vehicle throughput have the same color). Significantly more spectrum is needed when that spectrum is allocated in separate bands for vehicles and hotspots, when compared to all devices using shared spectrum. Therefore, it is possible to obtain the same performance for vehicles and hotspots using significantly less spectrum when it is shared, compared with vehicles and hotspots using separate spectrum. As for the differences among the sharing schemes, coexistence and relay cooperation require the same bandwidth to achieve given throughputs. However, backhaul cooperation requires less spectrum for some throughputs. To achieve vehicle throughput of about 60 Mbps/km^2 or less, backhaul cooperation requires the same bandwidth as the other schemes. To achieve vehicle throughput between 60 and 80 Mbps/km^2 backhaul cooperation requires up to 15% less bandwidth compared to coexistence and relay cooperation. Vehicle throughput between 80 and 85 Mbps/km^2 can be achieved with backhaul cooperation but not with any other scheme in this scenario.

Less shared spectrum is required for other scenarios as well, when compared with vehicles and hotspots using separate spectrum. Fig. 12 shows required spectrum in a scenario of 50 m separation between outdoor hotspots. The findings for this scenario are similar to those for the base case, even though 2/3 of vehicle throughput is lost with sharing. Fig. 12 shows that any sharing scheme requires significantly less spectrum than V2X and unlicensed devices using separate bands. Also, coexistence and relay cooperation require the same bandwidth to achieve given throughputs, while backhaul cooperation requires less spectrum than the other sharing schemes for most throughputs. However, the differences among the strategies are less for the scenario in Fig. 12 than for the base case. This is because in a scenario with more outdoor hotspots, there is more interference, and thus lower throughput, which results in less bandwidth savings when the ITS band is shared.

The effect of population density on required bandwidth for different sharing strategies is not obvious, because quantities of both V2X and unlicensed devices vary. For locations with lower population density the differences among sharing strategies (and from no sharing) are greatly affected by the presence of outdoor hotspots, as shown previously for more populated locations, while backhaul cooperation results in significant bandwidth savings. Fig. 13 shows the required bandwidth to achieve given throughputs for a higher population density, and we find that all sharing strategies require significantly less bandwidth than vehicles and unlicensed devices using separate bands. This is because of the increased number of both V2X and unlicensed devices compared to previous scenarios. Although this scenario results in more mutual interference, it also results in more data being transmitted (compared to lower population densities), thus increasing spectrum efficiency with sharing. Also, here is no significant difference among the bandwidths required with the several sharing schemes, which is consistent to the fact that all schemes produce similar throughputs for higher population densities as shown previously.

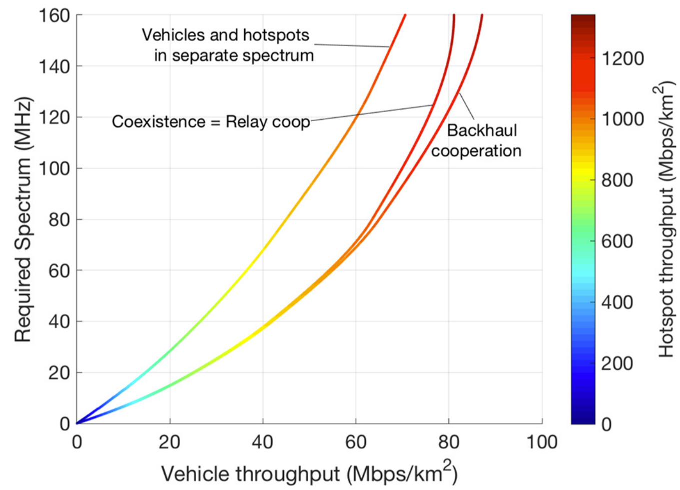

Fig. 11. Required spectrum to achieve given vehicular and hotspot throughputs, as a function of vehicular throughput. Points of equal color refer to equal hotspot throughput. Colors are coded in the bar (right). Assumptions are at base-case values (with “high” data rates).

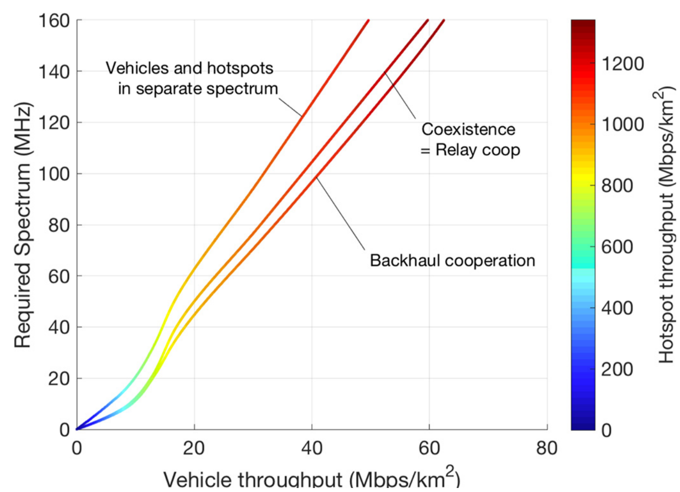

Fig. 12. Required spectrum to achieve given vehicular and hotspot throughputs, as a function of vehicular throughput, for different strategies. Distance between outdoor hotspots is 50 m, and the other assumptions are at base-case values (with “high” data rates).

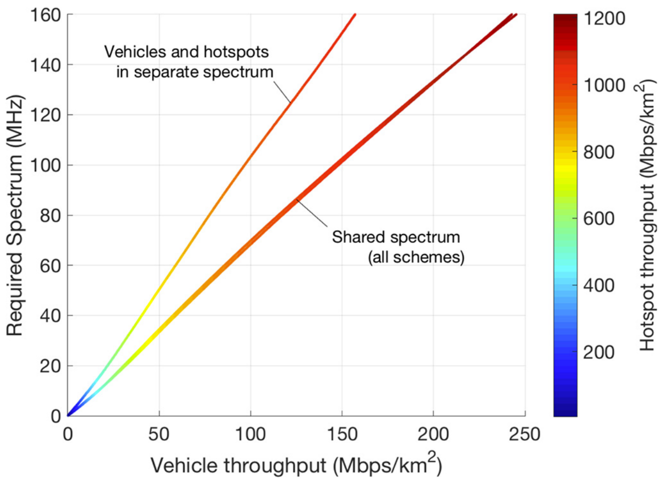

Fig. 13. Required spectrum to achieve given vehicular and hotspot throughputs, as a function of vehicular throughput, for different strategies. Population density is 4,000 people/km2, and the other assumptions are at base-case values (with "high" data rates).

V. Conclusion

We address the issues of how much spectrum should be available for ITS, whether that spectrum should be shared with unlicensed devices, as has been proposed by the FCC, and if so, how sharing should be implemented. We considered the scenario in which safety messages are transmitted over spectrum that is not shared, while V2X and unlicensed devices share spectrum on a co-equal basis to carry non-safety-critical information such as Internet traffic. (This is consistent with proposals such as [11].) We consider V2X and unlicensed devices operating in separate bands, and three possible sharing schemes that involve coexistence or cooperation. If spectrum is allocated exclusively to ITS, there are realistic scenarios where allocating spectrum far in excess of what is used for safety enhances social welfare, and there are also scenarios where too much spectrum has already been allocated. The bandwidth that maximizes social welfare is sensitive to uncertain factors such as the penetration of V2X, data rates, and the opportunity cost of 5.9 GHz spectrum. Because of this uncertainty, allocating spectrum exclusively runs the risk of not providing enough spectrum for welfare-enhancing ITS.

This uncertainty is less problematic if spectrum is shared. We also found that it is highly efficient to share spectrum allocated for ITS with unlicensed devices. Vehicles and unlicensed devices using separate bands might require significantly more bandwidth than is required to achieve the same throughputs in shared spectrum. This is true for scenarios that represent the relevant range of population densities, penetrations of vehicular devices and data rates of Internet traffic, and whether unlicensed devices are indoors or outdoors. While sharing is spectrally efficient when usage of V2X and unlicensed devices are predictable, it is even better in scenarios where data rates and/or penetration are much lower than expected due to the uncertainty discussed above, because even if an increase in ITS spectrum is not justified when spectrum is allocated exclusively, that spectrum is still well used by unlicensed devices when ITS spectrum is shared.

Backhaul cooperation can be more efficient than simpler coexistence. However, the magnitude of this advantage is scenario-dependent. Given that cooperation would require regulations that are far more complex than coexistence [10], it is unlikely that the benefits of cooperation outweigh the cost of implementing it. Moreover, the other cooperation scheme examined (relay cooperation) does not result significantly different from those of simpler coexistence. Therefore, a nationwide mandate for relay cooperation over coexistence would probably not be worth the technical and regulatory cost.

In the recent policy debate over ITS spectrum, it has generally been assumed that the size of the ITS band is fixed and the question is whether to share with unlicensed devices. In cases where the bandwidth available to vehicles is fixed, we found that the throughput achievable in shared spectrum can be lower than the throughput in exclusive spectrum. However, there is no reason why the bandwidth of the ITS band cannot be increased if we allow unlicensed devices to share the ITS band. If spectrum policymakers wish to give V2X better throughput than they could achieve in the existing ITS band after unlicensed devices are allowed to share, then regulations could increase the size of the ITS band while still giving unlicensed devices access, rather than prevent unlicensed devices form using the current ITS band. (Again, sharing the ITS band might exclude the portion of the ITS band reserved for safety messages.) Under these circumstances, vehicles and unlicensed devices would achieve the same throughput performance in shared spectrum while using less bandwidth overall. Besides, throughput to unlicensed devices in shared spectrum is not significantly lower than in exclusive spectrum. Therefore, sharing spectrum allocated for ITS with unlicensed devices effectively represents extra bandwidth for those devices, without compromising their throughput performance.

This work is based on the current U.S. DOT choice for V2X technology. However, the conclusions derived in this paper allow us to discuss a scenario in which the emerging cellular V2X technology (C-V2X) is used instead of DSRC. We have found that indoor hotspots do not significantly degrade DSRC throughput and vice-versa because of high path loss between where these devices are typically located. For the same reasons, C-V2X throughput might not significantly degrade as well, which would make sharing spectrum between C-V2X and hotspots efficient for the scenarios that we considered. However, there are enough differences to make this conclusion uncertain. First, C-V2X is claimed to have larger communications ranges than DSRC [38]. If true, this would both increase the number vehicles that can connect with an RSU, which increases benefit, and increase the interference that RSUs and vehicles can cause, which decreases benefit. Second, C-V2X does not employ listen-before-talk mechanisms used in DSRC and Wi-Fi, which may cause more degradation to outdoor hotspots than DSRC and make sharing less efficient. Further work is needed. However, other technologies (e.g., 5G) may enhance sharing efficiency. 5G may increase the capacity of cellular networks. While it is difficult to predict whether capacity will exceed traffic, if it does then there could be relatively less traffic over V2X, and sharing V2X spectrum with unlicensed devices may be even more efficient than we have estimated.

ACKNOWLEDGMENT

The authors thank Veniam Networks, STCP, the municipality of Porto, and the Future Cities Project (EU FP7).

REFERENCES

[1] J. Lansford, J. B. Kenney, P. Ecclesine, T. Yucek, and P. Spaanderman, "Final report of DSRC coexistence tiger team," IEEE, Piscataway, NJ, USA, Rep. IEEE 802.11-15/0347r0, 2015.
[2] *Report and Order 03-324*, U.S. Federal Commun. Commission, Washington, DC, USA, 2004.
[3] 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.
[4] "Revision of part 15 of the commission's rules to permit unlicensed national information infrastructure (U-NII) devices in the 5 GHz band. Notice of proposed rulemaking 13-22," U.S. Federal Commun. Commission, Washington, DC, USA, Rep. ET Docket 13-49, 2013.
[5] A. K. Ligo, J. M. Peha, P. Ferreira, and J. Barros, "Throughput and economics of DSRC-based Internet of Vehicles," *IEEE Access*, vol. 6, pp. 7276–7290, 2017.
[6] *Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016–2021*, Cisco, San Jose, CA, USA, 2017.
[7] A. K. Ligo and J. M. Peha, "Cost-effectiveness of sharing roadside infrastructure for Internet of Vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 7, pp. 2362–2372, Jul. 2018.
[8] 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," in *Proc. 43rd Telecommun. Policy Res. Conf.*, 2015, p. 39.
[9] A. K. Ligo, J. M. Peha, and J. Barros, "Throughput and cost-effectiveness of vehicular mesh networks for Internet access," in *Proc. IEEE Veh. Technol. Conf.*, Montreal, QC, Canada, 2016, pp. 1–7.
[10] J. M. Peha, "Sharing spectrum through spectrum policy reform and cognitive radio," *Proc. IEEE*, vol. 97, no. 4, pp. 708–719, Apr. 2009. [Online]. Available: www.ece.cmu.edu/∼peha/papers.html
[11] "Comments of Qualcomm incorporated," Qualcomm, San Diego, CA, USA, Rep. ET Docket 13-49, 2013.
[12] B. Cheng, H. Lu, A. Rostami, M. Gruteser, and J. B. Kenney, "Impact of 5.9 GHz spectrum sharing on DSRC performance," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Turin, Italy, 2017, pp. 215–222.
[13] J. Liu, G. Naik, and J.-M. J. Park, "Coexistence of DSRC and Wi-Fi: Impact on the performance of vehicular safety applications," in *Proc. IEEE Int. Conf. Commun.*, vol. 3. Paris, France, 2017, pp. 1–6.
[14] Y. Park and H. Kim, "On the coexistence of IEEE 802.11ac and WAVE in the 5.9 GHz band," *IEEE Commun. Mag.*, vol. 52, no. 6, pp. 162–168, Jun. 2014.
[15] G. Naik, J. Liu, and J.-M. J. Park, "Coexistence of dedicated short range communications (DSRC) and Wi-Fi: Implications to Wi-Fi performance," in *Proc. IEEE INFOCOM*, Atlanta, GA, USA, 2017, pp. 1–9.
[16] J. Wang, T. Wu, Y. Liu, W. Deng, and H. Oh, "Modeling and performance analysis of dynamic spectrum sharing between DSRC and Wi-Fi systems," *Wireless Commun. Mob. Comput.*, vol. 16, no. 16, pp. 2743–2758, 2016.
[17] H. J. Kim and J. M. Peha, "Detecting selfish behavior in a cooperative commons," in *Proc. 3rd IEEE Symp. New Front. Dyn. Spectr. Access Netw. (DYSPAN)*, Chicago, IL, USA, 2008, pp. 1–12. [Online]. Available: www.ece.cmu.edu/∼peha/papers.html
[18] K. N. Ramachandran, E. M. Belding, K. C. Almeroth, and M. M. Buddhikot, "Interference-aware channel assignment in multi-radio wireless mesh networks," in *Proc. 25th IEEE INFOCOM Int. Conf. Comput. Commun.*, Barcelona, Spain, 2006, pp. 1–12.
[19] T. W. Hazlett and M. Honig, "Valuing spectrum allocations," *Mich. Telecomm. Tech. Law Rev.*, vol. 23, p. 45, May 2016.
[20] *Planning for the Future of Transportation: Connected Vehicles and ITS*, U.S. Dept. Transp., Washington, DC, USA, 2015.
[21] A. Aittokallio, "U.S. AWS-3 auction raises record $45 billion," Telecoms.com, London, U.K., Accessed: Jan. 30, 2015. [Online]. Available: http://telecoms.com/395312/us-aws-3-auction-raises-record-45-billion/
[22] J. M. Peha, "Cellular competition and the weighted spectrum screen," in *Proc. TPRC41 Res. Conf. Commun. Inf. Internet Security*, 2013, pp. 1–23.
[23] M. Alotaibi, M. A. Sirbu, and J. M. Peha, "Impact of spectrum aggregation technology and spectrum allocation on cellular network performance," in *Proc. IEEE Conf. Dyn. Spectr. Access Netw. (DySPAN)*, Stockholm, Sweden, 2015, pp. 326–335.
[24] R. Katz, "Assessment of the economic value of unlicensed spectrum in the United States," in *Proc. TPRC 42nd Res. Conf. Commun. Inf. Internet Policy*, 2014, p. 99.
[25] 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. Chichester, U.K.: Wiley, 2009.
[26] *WRV210 Wireless-G VPN Router: RangeBooster*, Cisco, San Jose, CA, USA, 2014. [Online]. Available: https://www.cisco.com/c/en/us/products/collateral/routers/wrv210-wireless-g-vpn-router-rangebooster/data_sheet_c78-502735.html.
[27] *Federal Motor Vehicle Safety Standards; V2V Communications—Notice of Proposed Rulemaking (NPRM)*, U.S. Dept. Transp., Washington, DC, USA, 2016.
[28] *Cisco Visual Networking Index (VNI): North America*, Cisco, San Jose, CA, USA, 2016.
[29] J. Ozer. (2016). *How Netflix Pioneered Per-Title Video Encoding Optimization—Streaming Media Magazine*. [Online]. Available: http://www.streamingmedia.com/Articles/Editorial/Featured-Articles/How-Netflix-Pioneered-Per-Title-Video-Encoding-Optimization-108547.aspx
[30] J. Engebretson. (2016). *iGR: Average Monthly Broadband Usage Is 190 Gigabytes Monthly Per Household—Telecompetitor*. [Online]. Available: http://www.telecompetitor.com/igr-average-monthly-broadband-usage-is-190-gigabytes-monthly-per-household/
[31] R. Hallahan and J. M. Peha, "The business case of a network that serves both public safety and commercial subscribers," *Telecommun. Policy*, vol. 35, no. 3, pp. 250–268, 2011. [Online]. Available: www.ece.cmu.edu/∼peha/papers.html
[32] *Outdoor Mobility Design Guide*, Cisco, San Jose, CA, USA, 2010. [Online]. Available: https://www.cisco.com/c/en/us/support/docs/wireless/4400-series-wireless-lan-controllers/111902-outdoor-mobi-guide-00.html
[33] R. Hallahan and J. M. Peha, "Quantifying the costs of a nationwide public safety wireless network," *Telecommun. Policy*, vol. 34, no. 4, pp. 200–220, 2010. [Online]. Available: www.ece.cmu.edu/∼peha/papers.html
[34] J. M. Peha and F. A. Tobagi, "Cost-based scheduling and dropping algorithms to support integrated services," *IEEE Trans. Commun.*, vol. 44, no. 2, pp. 192–202, Feb. 1996.
[35] G. Naik, J. Liu, and J.-M. J. Park, "Coexistence of wireless technologies in the 5 GHz bands: A survey of existing solutions and a roadmap for future research," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 3, pp. 1777–1798, 3rd Quart., 2018.
[36] J. Lansford, J. B. Kenney, and P. Ecclesine, "Coexistence of unlicensed devices with DSRC systems in the 5.9 GHz ITS band," in *Proc. IEEE Veh. Netw. Conf.*, Boston, MA, USA, 2013, pp. 9–16.
[37] K.-H. Chang, "Wireless communications for vehicular safety," *IEEE Wireless Commun.*, vol. 22, no. 1, pp. 6–7, Feb. 2015.
[38] A. K. Ligo, *Connected Vehicles for Internet Access: Deployment and Spectrum Policies*, Carnegie Mellon Univ., Pittsburgh, PA, USA, and Univ. Porto, Porto, Portugal, 2018.

Alexandre K. Ligo (M'11) received the B.S. degree in electrical engineering from the University of São Paulo, Brazil, and the dual Ph.D. degrees in engineering and public policy from Carnegie Mellon University (CMU), USA, and the University of Porto, Portugal. He is a Post-Doctoral Researcher with CMU. His research interests include Internet policy, connected and automated vehicles, wireless communications, spectrum management, and benefit-cost analysis. Before joining academia, he worked over 15 years in research, engineering, and consulting in the ICT industry. He is a member of the Institute of Telecommunications, Porto.

Jon M. Peha (F'12) received the B.S. degree from Brown University and the Ph.D. degree in electrical engineering from Stanford University. He is a Full Professor with CMU with experience in industry, government, and academia. In government, he served at the Federal Communications Commission as a Chief Technologist, in the White House as an Assistant Director of the Office of Science and Technology, in the House Energy and Commerce Committee, he was responsible for telecom and e-commerce issues, and at USAID he helped to launch and led a U.S. Government interagency program to assist developing countries with information infrastructure. In industry, he has been a Chief Technical Officer for three high-tech companies. He is a Professor with the Department of Engineering and Public Policy and the Department of Electrical and Computer Engineering, CMU, and a Former Associate Director of the university's Center for Wireless and Broadband Networking. His research spans technical and policy issues of information networks, including spectrum management, broadband, communications for emergency response, vehicular networks, smart cities, privacy, and cybersecurity. He was a recipient of the FCC's "Excellence in Engineering Award," the IEEE ComSoc's TCCN Publication Award for career contributions, and the Brown Engineering Medal. He is a fellow of AAAS and was selected as one of 40 Featured AAAS Science and Technology Policy Fellows of the last 40 years ("40@40").