Connected Vehicles for Internet Access: Deployment and Spectrum Policies

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Engineering and Public Policy

Alexandre K. Ligo

B.S., Electrical Engineering, University of São Paulo
M.S., Business Administration, Getúlio Vargas Foundation

Carnegie Mellon University
Pittsburgh, PA

December, 2018

ACKNOWLEDGEMENTS

Thank God I am done!

I thank my wife Thais, because this work would never be completed without her positive attitude and encouragement, especially during four overseas moves. She always kept her mind straight when I lost mine. I thank my son Alexandre Ligo Filho, who at 9 years old is my model of happiness and hard-work. He is my greatest accomplishment ever. I am grateful to my mother, sister, parents-in-law, brother-in-law, uncle Gilmar, and all the family.

I thank Professor Jon Peha, my advisor and chair of the doctoral committee, for his brilliant insights and endless intellectual and material support. He was always willing to thoroughly revise all my work, even when I was dangerously close to the deadlines. I also thank my co-advisor Professor João Barros for the contributions and support for this research, and the other members of the committee who generously gave many comments and suggestions for this work: Professor Douglas Sicker, Professor Tim Brown, Professor João Claro, and Professor José Silva Matos.

I am grateful to the teams at Instituto de Telecomunicações Porto, Veniam Networks, the municipality of Porto and STCP for providing the data used in this work. I also thank Professor Ana Aguiar and her team for hosting me in her lab during my research in Portugal, and Professor Pedro Ferreira at CMU.

I express my gratitude to the several institutions that provided financial support to make this work possible. This research was supported by the CMU-Portugal Partnership (scholarship no. SFRH/BD/51564/2011), by the Portugal Foundation for the Science and Technology (ref.

UID/EEA/50008/2013), by the Future Cities Project (European Commission EU FP7 under grant no. 316296), 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 by the National Science Foundation (grant no. 1547237).

ABSTRACT

Internet traffic from mobile users has been growing sharply. To meet the needs of those users, it is important to expand capacity of networks that provide Internet access in cost-effective ways. This capacity has traditionally been provided by cellular networks. However, expanding the capacity of those networks alone may not be the most cost-effective way to meet the present and future growth of mobile Internet under some circumstances. In this dissertation, we show that networks of connected vehicles can be an important way to complement the capacity of cellular networks to provide mobile Internet access under several scenarios. Connected vehicles may soon be widely deployed, forming mesh networks of short-range connections among vehicles and between vehicles and roadside infrastructure. These connections are collectively referred to as vehicle-to-everything, or V2X. Deployment of connected vehicles and infrastructure is primarily intended to enhance road safety, and the U.S. Department of Transportation has recently proposed a mandate of V2X devices in vehicles using Dedicated Short Range Communications (DSRC) technology. Other applications are also envisioned that include Internet access in vehicles connecting to roadside infrastructure serving as gateways to the Internet.

In this work, we find that V2X-based networks are more cost-effective than cellular to provide Internet access, in scenarios which DSRC devices are mandated in vehicles to enhance road safety. This is true initially for densely populated urban areas, but over time V2X-based networks would be cost-effective in less populated areas as well, as long as Internet traffic or penetration of V2X devices grow as expected.

Local and state governments are expected to deploy roadside infrastructure for safety applications. If that infrastructure is shared with Internet Service Providers for a fee, then V2X-

based networks are cost-effective in locations with even lower population densities than the locations where it is cost-effective to deploy infrastructure for Internet access only. Moreover, the sharing fee could help governments save in infrastructure costs. We find the pricing strategies that maximize either cost-effectiveness or government savings. We estimate that governments could save about one-fifth of the total cost to deploy safety infrastructure nationwide in the U.S., if fees are set to maximize government savings. Although we find that these prices may differ from the pricing strategy that maximizes cost-effectiveness, maximizing government savings results in near-optimal cost-effectiveness.

The U.S. Federal Communications Commission has allocated 75 MHz of spectrum to be used exclusively by DSRC devices, and it has been hotly debated whether all or part of that bandwidth should be shared with unlicensed devices. We find that it is highly efficient to share any spectrum allocated to V2X communications beyond the portion of that spectrum that is needed for safety-critical DSRC messages. V2X and unlicensed devices require up to 50% less bandwidth on shared spectrum to achieve given throughputs, compared to V2X and unlicensed devices using separate bands. We conclude that the spectrum available for V2X should be maintained or increased, as long as much of that spectrum is shared with non-V2X devices.

Conclusions are derived from an engineering-economic approach, in which part of the assumptions are based on data from a citywide deployment of connected vehicles in Portugal. The data is used in a detailed and realistic packet-level simulation model of V2X-based networks used to provide Internet access with DSRC technology. In some scenarios, the simulation also includes unlicensed devices using Wi-Fi technology. The results of the network simulation are then fed into engineering-economic models to compare costs of V2X-based networks with costs of macrocellular networks to carry given amounts of Internet traffic, and to estimate other measures such as government revenues and spectrum usage. Those measures help inform decisions about where and when to deploy V2X-based networks, decisions about whether and how to promote public-private partnerships to deploy V2X infrastructure, and decisions about sharing spectrum used for V2X communications with non-V2X devices.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ACRONYMS

3GPP	3rd Generation Partnership Project
5GAA	5G Automotive Association
C-V2X	Cellular vehicle-to-everything
DOT	U.S. Department of Transportation
DSRC	Dedicated short range communications
FCC	U.S. Federal Communications Commission
IP	Internet protocol
ISP	Internet Service Provider
ITS	Intelligent Transportation Systems
LTE	Long Term Evolution
NPRM	Notice of Proposed Rulemaking
NPV	Net Present Value
OBU	Onboard unit
RSU	Roadside unit
V2I	Vehicle-to-infrastructure
V2V	Vehicle-to-vehicle
V2X	Vehicle-to-everything

Chapter 1. Introduction

1.1. *Scope of this work*

Vehicular communications can be used for several purposes, which range from improving road safety to providing Internet access. One way to enable vehicular communications is by deploying vehicular mesh networks. In such networks, vehicles directly send packets to each other over short distances, in what is known as vehicle-to-vehicle or V2V communications. Vehicles can also communicate over short distances with fixed access points placed near roads, in what is called vehicle-to-infrastructure or V2I communications. V2V and V2I are collectively referred as vehicle-to-everything or V2X communications[1]. Routers in cars and in roadside infrastructure may form a mesh network of V2V and V2I links to forward data between vehicles and endpoints located outside the V2X mesh, such as central offices or the Internet.

This work is about the use of V2X-based mesh networks as a way to provide Internet access. Internet traffic has grown sharply and probably will continue to increase, especially from mobile users. Mobile traffic has traditionally been carried over cellular networks. If part of the Internet traffic may be carried through alternatives such as V2X-based networks, then these networks could complement the capacity of cellular networks to meet the demand growth.

[1] Some authors use V2X to include other forms of communications such as vehicle-to-pedestrian, bicycle, network, etc. However, we use V2X to refer to V2V and V2I because the issues addressed in this work are more closely related to communications among vehicles and between vehicles and roadside infrastructure.

We inform decisions regarding whether V2X-based networks should be deployed for Internet access or not, as well as where, when and how to deploy them. To inform those decisions we address two interrelated issues. First, under what conditions are V2X-based networks more cost-effective than cellular networks for carrying Internet traffic? By determining those conditions with respect to population density, penetration of V2X devices in vehicles, data rates of Internet traffic, costs of equipment and spectrum, and characteristics of the cellular network, we are able to inform where and when V2X-based networks are cost-effective and thus should be deployed. The second issue is what are the best design choices for the V2X network for each set of conditions. These choices include quantities of infrastructure and spectrum, which are chosen as the most cost-effective under different conditions. Other choices that we examine are whether infrastructure and spectrum should be deployed exclusively for in-vehicle Internet access, or rather those resources should be shared for other uses.

1.2. V2X communications

Connected vehicles may soon be widely deployed as an important element of Intelligent Transportation Systems (ITS). Standardized technology now exists that would support vehicular communications, which include mesh-based vehicular networks running Internet protocols such as IP. In a vehicular mesh network, routers are placed in onboard units (OBUs) in automobiles. Routers are also placed in fixed infrastructure of roadside units (RSUs) deployed for communications purposes. OBUs and RSUs thus enable V2V and V2Icommunications.

Since 1999, the U.S. Federal Communications Commission (FCC) has allocated 75 MHz of spectrum in the 5.9 GHz band for use in vehicular communications (U.S. Federal Communications Commission 1999). (Besides, 70 MHz has been allocated for vehicular communications in Europe, and similar bandwidths have been allocated or planned for allocation in Japan and other regions.) This spectrum is known as the *ITS band*. Currently, the leading technology that enables V2X-based mesh networks is called *Dedicated Short Range Communications*, or *DSRC* (Kenney 2011; Campolo and Molinaro 2013). This technology enables transmissions over short distances of up to 1,000 meters in open space, and up to about 300 meters in urban areas. In 2004, the FCC released a Report and Order that determines that only OBUs and RSUs based on DSRC technology are allowed to operate in the ITS band (U.S. Federal Communications Commission 2004)[2]. In 2016, the U.S. Department of Transportation (DOT) has proposed rulemaking in 2016 to mandate DSRC devices to be deployed, at first in all new light vehicles that include passenger cars and certain types of trucks[3]. In the proposed rule it is expected that all new vehicles of those types are deployed with DSRC devices in no more than five years after the rule is adopted (U.S. Department of Transportation 2015, 2016).

The spectrum allocation and proposed rulemaking suggest that V2X technology may be widely deployed. The primary purpose for that deployment is to enhance automotive safety, and extensive research has been done on the use of V2X communications for safety applications (Kenney 2011; U.S. Department of Transportation 2015; Mecklenbrauker et al. 2011). One example is forward collision warning, which is the exchange of warning messages between vehicles *en route* of collision. A vehicle may broadcast V2V messages with its position, speed, heading and other information. Other vehicles in the surroundings receive those messages, and

---

[2] According to the FCC's R&O 03-324 of 2004, vehicular communications based on DSRC, military radar and fixed satellite uplink operations are co-primary in the 5.9 GHz band.

[3] In the U.S. DOT proposed rule, light vehicles are defined as "passenger cars, vans, minivans, sport utility vehicles, crossover utility vehicles and light pickup trucks with a gross vehicle weight rating (GVWR) less than or equal to 10,000 pounds" (U.S. Department of Transportation 2016).

it they determine a dangerous condition, may either warn the drivers or take autonomous actions such as changing lanes or reducing speed. This example illustrates the use of V2V communications for low-latency, safety-critical applications. With direct V2V communications, vehicles avoid exchanging messages with remote servers, which might increase latency. Another example is of roadside communications infrastructure connected to a transportation authority's facility. Upon receiving information about a dangerous weather condition at a certain locality, the transportation authority has the information broadcast by units of roadside infrastructure within the given locality. The V2I messages are received by cars passing nearby, which act according to the hazard. Requirements of safety-critical communications are relatively challenging with respect to reliability and latency. (On the other hand, several safety applications are not data-intensive, so high throughput is not a concern.)

The U.S. DOT and several other organizations, including the 3rd Generation Partnership Project (3GPP) and the 5G Automotive Association (5GAA) have envisioned several applications that involve V2X communications (Campolo et al. 2017; J. Chang et al. 2015). Those applications include both use cases that are directly related to enhancing road safety, and others in which V2X gives each vehicle the ability to collect, disseminate, and receive information about the vehicle surroundings, and gives the vehicle and its occupants the ability to interact more fully on the Internet. We refer to use cases not directly related to road safety-critical applications as *non-safety*. Examples of non-safety applications include vehicle and road-related services such as navigation and centralized traffic congestion. Upon collecting V2I information about the load of vehicles in major roads at peak hours, the transportation authority has information about the best available routes broadcast in real-time by roadside infrastructure. The V2I messages are received by passing cars that then choose the faster routes while avoiding congested one. Other examples of the use of non-safety applications based on V2X-based communications range from toll collection to sensor data gathering by automated vehicles, as well as business and entertainment activities engaged by vehicle occupants ranging such as voice calls, online video and online gaming. Several of those applications depend on in-car Internet access, which some have named as the “Internet of Vehicles”.

DSRC standards allow part of the spectrum allocated in the ITS band to be used for non-safety applications (Zeadally et al. 2010; Uzcategui and Acosta-Marum 2009; Campolo and Molinaro 2013). In this case, V2X safety messages will either be sent over exclusive channels where no other type of messages is allowed, or will have higher priority than non-safety messages in channels where the latter are allowed (Kenney 2011).

1.3. Problem statement: how to expand network capacity for mobile Internet?

Internet traffic has increased dramatically over time. Part of that growth is associated with mobile users. Mobile Internet traffic has grown 18-fold in the past 5 years, and will likely continue to grow that way or even faster (CIsco 2017). The 3GPP predicts new forms of mobile traffic will soon emerge in the vehicular context (Campolo et al. 2017), which include video and other types of data from sensors, as well as in-vehicle Internet access for occupants in self-driving cars.

It is debated whether and how the capacity of today’s access networks to carry Internet traffic is going to keep up with the dramatic growth in demand. Internet traffic from mobile users has been served primarily by cellular networks. However, expanding the capacity of those networks alone may not be the most cost-effective way to meet the present and future growth of mobile Internet under some circumstances.

1.4. V2X-based networks to complement cellular capacity

If part of mobile Internet traffic could be carried over networks other than cellular, then the growth in demand might be met with less cellular infrastructure. Users connect to the Internet via several wired and wireless access technologies, which include cellular, cable, optical fiber, Wi-Fi, etc. Macrocellular networks are typically used to provide mobile Internet service for fast-moving users. Instead of serving all Internet traffic through cell towers, Internet Service Providers (ISPs) often use alternative infrastructure, such as Wi-Fi hotspots connected to the Internet via cable or fiber, to deviate some Internet traffic from cell towers. By doing this, ISPs avoids using expensive cellular infrastructure to carry whatever data that can be served by cheaper fixed infrastructure. This is often referred to as mobile data offload over Wi-Fi, which typically works well with slow-moving users who access the Internet with their laptops or smartphones connected to Wi-Fi hotspots. However, Wi-Fi hotspots are not suitable to offload data of fast-moving users such as those in vehicles, mainly because of the time required for a user to associate and authenticate to a hotspot.

This work is about using V2X-based mesh networks using DSRC technology, as a novel way to provide mobile Internet access to devices in automobiles, as well as devices carried by pedestrians or placed in locations near roads. If the cost to carry some part of traffic over V2X-based networks is lower than the cost over cellular infrastructure, then V2X devices and spectrum could be deployed to offload that part of traffic at a lower cost than expanding cell towers alone. In this case, V2X-based networks are more cost-effective than expanding cellular networks alone. While work such as (Campolo et al. 2017) has predicted the need of V2X communications to offload vehicular traffic from cellular networks, one of our contributions is to quantify benefits of such offloading with respect to economic and technical measures.

We determine the conditions in which V2X networks are more cost-effective than cellular to provide Internet access. One factor to consider is the population density of a location, which is correlated with the density of vehicles there. If the density of vehicles equipped with V2X devices is proportional to the volume of Internet traffic demanded by vehicular users in a location, then there is a relationship between population density and the demand for Internet traffic. For each of several population densities, we examine whether the deployment of V2X devices and spectrum to offload cellular traffic is more cost-effective than deploying cellular networks alone. By doing that, we determine the range of population densities where V2X-based networks are cost-effective. Other conditions that we examine include penetration of V2X devices in vehicles, data rates of Internet traffic per vehicle, costs of equipment and spectrum, and characteristics of the cellular network.

Another condition that may affect cost-effectiveness of V2X-based networks for Internet access is whether V2X resources are deployed for other purposes. Spectrum has been allocated for V2X communications, and it is possible that V2X devices are mandated for safety purposes. In this case, some of the costs to deploy V2X-based networks will be incurred to enhance road safety, regardless of whether those networks are used for other purposes. If those costs are incurred anyway, then the only V2X costs that matter for offload are the incremental costs necessary to provide Internet access, such as Internet-connected RSUs. Another of our contributions is to determine if these incremental costs are less than expanding cellular infrastructure.

The conditions in which V2X-based networks are cost-effective are interrelated with design choices, i.e. how those networks are deployed. One of those choices is the amount of infrastructure to deploy. We examine if and how the quantity of RSUs to deploy in a location depends on population density and other conditions. With that analysis we determine what is the quantity of RSUs that is most cost-effective for given conditions. Other design choices that we examine in this work include whether RSUs are deployed for Internet alone or rather shared with other purposes (e.g. V2I communications for safety applications), the amount of ITS spectrum to allocate for V2X communications, and whether the ITS band should be shared with devices that do not use V2X communications (e.g. unlicensed devices as proposed in recent FCC proceedings).

1.5. Implications of alternative V2X technologies

The findings in this work were derived from V2X networks based on DSRC. This technology has been developed for vehicular communications for about 20 years, by entities in the transportation, communications and automotive sectors. However, more recently a technology known as cellular V2X (C-V2X) has emerged, and in this dissertation, we discuss how our findings apply if vehicular networks are deployed with C-V2X rather than DSRC.

C-V2X technology might be either an alternative to DSRC in greenfield deployments, or a complement if DSRC has already been deployed (Qualcomm 2018). C-V2X is currently being incorporated in the latest and future releases of 3GPP cellular standards. In 2016, the first version of C-V2X was released (3GPP Rel. 14) to support the requirements of vehicular safety applications through cellular networks (Flore 2016). This includes the specification of a new cellular interface, which is known as sidelink or PC5. The sidelink is a V2V interface where devices send messages directly to each other, which will work even without coverage from cellular infrastructure, using the ITS band. The sidelink interface is intended to meet safety latency requirements while being independent from cellular coverage. In the first version, this interface was optimized for safety applications where reliability and latency requirements must be met, but throughput capacity is of lesser concern. An example is forward collision warning, which is based on short messages being transmitted every 100 milliseconds. The newest version of the specification (3GPP Rel. 15) supports high throughput applications, which include high definition maps and live video over the sidelink (Qualcomm 2018; ETSI and 3GPP 2018). Moreover, the allocation of the ITS band for the sidelink helps ensure that vehicles subscribed to different cellular operators (or no operator at all) can exchange messages via the sidelink (ETSI 2018).

On the other hand, V2I communications in C-V2X can be based either on the new sidelink PC5, or on the traditional LTE Uu interface (ETSI 2018). When using the sidelink PC5 interface, the C-V2X RSU is equipped the same way as a vehicular device, with the difference that the RSU is static and possibly connected to the cellular network. In this case, the V2I link operates in the ITS band. On the other hand, a C-V2X RSU may alternatively use the traditional LTE Uu interface, which connects mobile users and the cell tower. The Uu interface operates on LTE bands around 2 GHz and below, and possibly on future 5G bands licensed to cellular operators (Molina-Masegosa and Gozalvez 2017; Campolo et al. 2017). In this form of V2I communications, if the vehicle is subscribed to an operator other than the one that owns the RSU, then the roaming mechanisms of LTE shall be used.

To date, there has been no consensus about whether DSRC, C-V2X or both should be deployed for vehicular communications. Proponents of DSRC typically include transportation agencies and car manufacturers. They argue that DSRC is a well-known technology based on IEEE 802.11 standards, which has been successfully tested for several years, and for that reason its capacity and reliability has been proven. Moreover, in the U.S. the ITS band is currently regulated for DSRC use (U.S. Federal Communications Commission 2004). On the other hand, proponents of C-V2X include telecommunications service providers and equipment vendors. They argue that DSRC is an aged technology that do not scale well with very high densities of vehicles, while C-V2X incorporates the latest radio technical features that are present in current and future cellular technologies. While the analysis and findings in this work are based on V2X networks using DSRC technology, we discuss whether and which of those findings apply to a scenario where V2X-based networks are deployed with C-V2X technology rather than DSRC.

1.6. Contents of the dissertation

The research questions and findings that are addressed in this dissertation are grouped in three parts. In the first part, the research questions are related to the conditions in which V2X-based networks are more cost-effective than cellular networks to provide Internet access. In this part, we consider specific design choices that represent infrastructure and spectrum being used exclusively by V2X devices for Internet access. It is considered that RSUs are deployed exclusively to provide Internet access. Besides, it is assumed that spectrum is allocated as currently defined in the ITS band by U.S. regulations. That is, spectrum is exclusively used by V2X devices, and a fixed amount of that spectrum used for non-safety applications, which include Internet access. These design choices have implications in performance and/or costs of V2X-based networks and are further varied in subsequent parts of the dissertation.

In the second part of our work we consider the possibility of governments sharing infrastructure with ISPs for a fee. It is likely that government agencies will deploy their own infrastructure for a variety of purposes. Local or state transportation agencies may deploy RSUs for V2X communications to support safety applications. Besides, other government agencies may deploy infrastructure for public lighting, surveillance or other purposes. That infrastructure would likely be deployed with poles or other supporting structure, and in places with access to power, and in some cases, with communications capability to reach remote offices of those agencies, or even the Internet. (In this work we call that “smart” infrastructure, since it might be part of smart city infrastructure prepared for future services that require communications capability.) If those governments share their infrastructure with ISPs, then the respective costs can also be shared, thereby reducing costs for both. That is possible if part of the infrastructure is made available for multiple purposes. For example, a transportation agency might make RSU deployed for safety available also for ISPs, in locations where those RSUs have Internet access with adequate bandwidth. (If bandwidth is less than required by the ISP, it could possibly be upgraded with a cost lower than deploying a new RSU with a pole, access to power and Internet, etc.) We consider the context of governments sharing either safety RSUs or smart streetlight infrastructure with ISPs for a fee.

In the second part of the dissertation we address three research questions. The first is whether ISPs sharing RSUs deployed by governments is more cost-effective than ISPs deploying their own infrastructure. The second research question is about the optimal pricing strategies that governments should adopt for sharing their infrastructure. If prices are too low, it may not be worthwhile for the government to share. If prices are too high, ISPs might choose to deploy their own RSUs, or not deploy at all. The third question is how much governments could save in infrastructure investment by sharing. We address this question by estimating government revenues from RSU sharing, which depend on the pricing strategies adopted. To address all those questions, we consider that a given amount of spectrum is allocated exclusively for vehicles.

In the third part we investigate how much spectrum should be available for ITS to maximize the cost-effectiveness of V2X-based Internet. We also examine whether part of that spectrum should be shared with unlicensed devices, as has been considered by the FCC, and if so, how such sharing should be implemented.

The results for each scenario of conditions and design choices are estimated using an approach that has three major parts. First, we extensively collected data from more than 800 vehicles in Portugal. Part of those vehicles formed a real large vehicular network deployed to offload data from cellular[4]. The second part is a detailed network simulation model. We have developed packet-level simulation software that uses the positions of vehicles from Portugal, to simulate a mesh network comprised by DSRC onboard units (OBU) in vehicles that connect to fixed roadside units (RSU) to gain access to the Internet. We use this simulation to estimate the throughput rate of Internet data that can be carried through the V2X-based network, and we assess how throughput is influenced by several conditions, ranging from bandwidth, densities of devices, and data rates of incoming traffic. The third part of the approach includes engineering-economic models that estimate costs and revenues as a function of throughputs. To do those estimates, we assume that mobile devices can use either cellular services or V2X-based networks. In a capacity-limited cellular network, a reduction of data from mobile devices that must be carried in the busy hour allows each cell tower to provide adequate capacity over a larger area, thereby reducing the number of costly towers that a cellular operator needs to cover a given region. We estimate cost savings from reducing the number of towers that would otherwise be required if V2X-based networks are not used for Internet access. Those savings are compared with the costs of deploying V2X-based networks, which may include costs of RSU infrastructure, spectrum and OBUs. The comparison is performed under a variety of relevant scenarios defined by several assumptions. These assumptions include the quantities and characteristics of infrastructure and spectrum, quantities of vehicles and other devices, data rates of incoming traffic, and unit costs of V2X and cellular components. The model is also used

[4] Data from the vehicular network deployed in Porto, Portugal, were made available for research at the University of Porto, which I am affiliated with.

to compute other results, such as government savings when examining sharing of infrastructure, and spectrum efficiency when addressing sharing of the ITS band.

This dissertation is organized as follows. In Chapter 2, we describe the portion of our engineering-economic approach that is common to all three parts of our work. This description includes the Portugal dataset, simulation model, economic analysis and base numerical assumptions. In Chapter 3 we address the first part of our work. We determine the conditions in which V2X networks are cost-effective in the context of in-vehicle devices are mandated, spectrum is allocated with fixed bandwidth, and Internet Service Providers (ISP) deploy V2X-based infrastructure for Internet access only. Then, in Chapter 4, we consider the second part of our work, which represents the context of governments sharing either safety RSUs or smart streetlight infrastructure with ISPs for a fee. In Chapter 5, we address the third part. We examine how much spectrum should be available for ITS to maximize the cost-effectiveness of V2X-based Internet, whether part of that spectrum should be shared with unlicensed devices, and if so, how spectrum should be shared.

Chapter 6 concludes the dissertation with an overall discussion about the benefits of using V2X-based networks to provide Internet access. That discussion includes an overview of our findings with respect to cost-effectiveness and other measures, in the several scenarios presented in the previous chapters. We also discuss whether and which findings of this work may apply to scenarios other than those considered in this work. Chapter 7 includes an outline of directions for future work.

Chapter 2. Baseline Method and System Model

To address the issues introduced in Chapter 1, in the chapters that follow we evaluate throughput, benefits, and costs of Internet access though V2X-based mesh networks under different conditions. This chapter describes the engineering-economic approach for the scenario of a mandate of V2X devices with exclusive use of ITS spectrum, and RSU infrastructure deployed for Internet access. We also describe the Portugal data and the core assumptions that are common to all subsequent analyses. The methodology is further extended for the analysis of infrastructure, spectrum and cooperation issues. Those extensions are described in the subsequent chapters.

We consider a heterogeneous scenario where mobile Internet traffic is carried over both V2X and macrocellular networks as represented in Figure 2.1. In this scenario, OBU-equipped vehicles are capable of connecting to the Internet in two ways. One is through macrocells. The other way is through V2X-based mesh networks comprised of multihop paths. These paths are formed by vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) links. Those links connect vehicles to RSUs that serve as gateways to the Internet. Whenever vehicles are able to establish a path to an RSU, they will use that path to send Internet traffic. Otherwise, the cellular network is used, which is assumed in this model as being ubiquitous. Safety messages and Internet packets are sent over separate channels in the ITS band. Therefore, in this model there is no mutual interference between safety and Internet traffic.

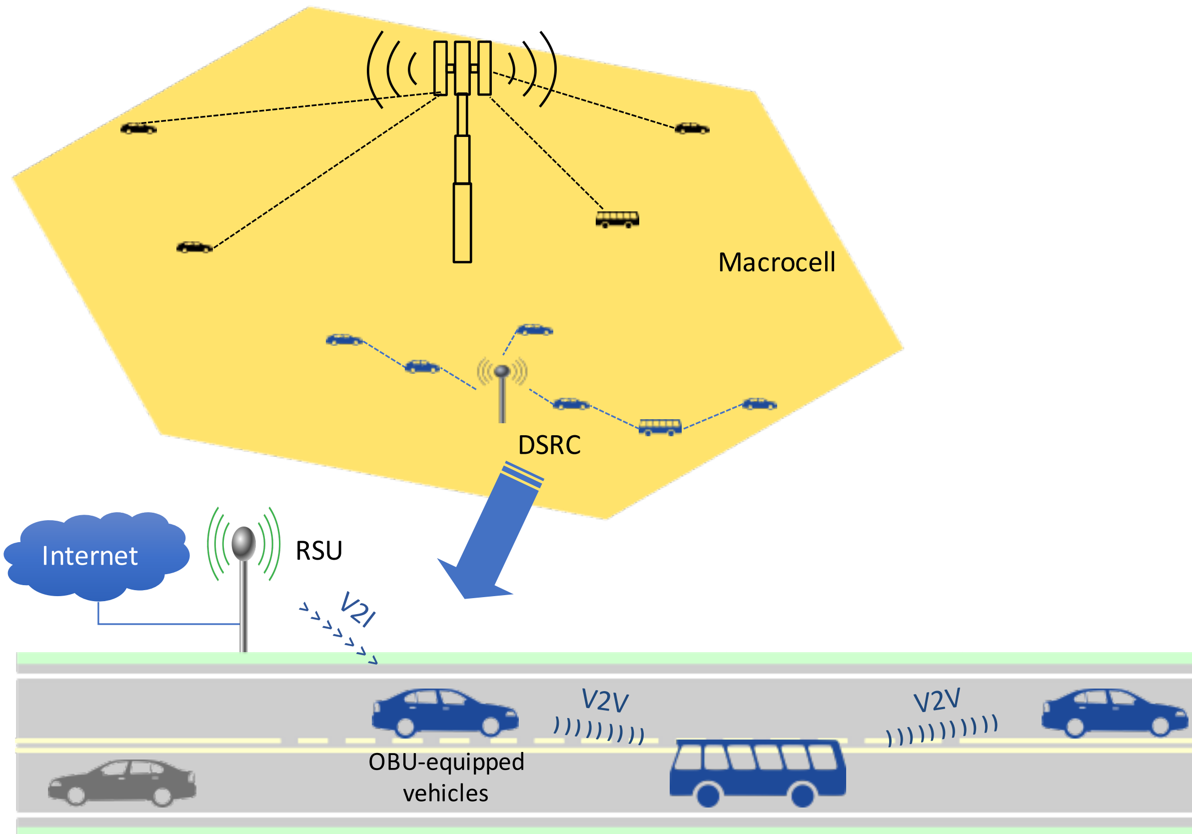

(a) Vehicular Internet access through heterogeneous networks (cellular and DSRC).

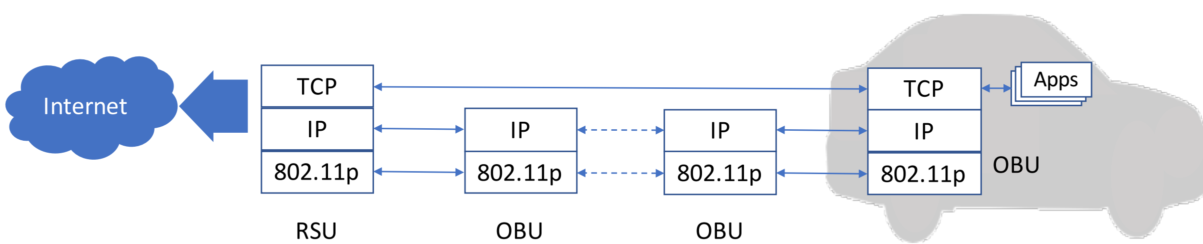

(b) Model of DSRC-based connection between an RSU and a vehicle equipped with an OBU.

Figure 2.1. Model of Internet access based over vehicular and microcellular networks.

To address research questions about the cost-effectiveness of vehicular communications for Internet access, we combine a two-step methodology with data collected from an actual vehicular network operating in Porto, Portugal. The methodology is used to estimate network throughput and costs and is outlined in Figure 2.2. The first step is to estimate how much vehicular Internet traffic, which would otherwise be carried by macrocells, can instead be carried by a V2X-based network under different conditions. To achieve this, we developed a detailed, packet-level simulation model to estimate the rate at which data is transferred between OBU-equipped vehicles and RSUs through single or multiple hops. That model makes use of representations of the elements that most influence the throughput rates, including the locations of RSUs and vehicles, the signal loss between them, the DSRC protocol itself, and multiple vehicles and RSUs exchanging data simultaneously in the same area.

Some of the realism of the simulation comes from measurement data collected from the city-wide vehicular network in Porto. For the analysis of the base-case scenario of a mandate and infrastructure and spectrum not being shared, Porto data is used in two ways. First, GPS positions from more than 800 buses and taxis in Porto are used to determine the positions of the vehicles in the simulation. Second, the received signal measured in the buses is used to verify whether the simulated signal loss (which influences transmission ranges and interference) is compatible with measured loss, on average.

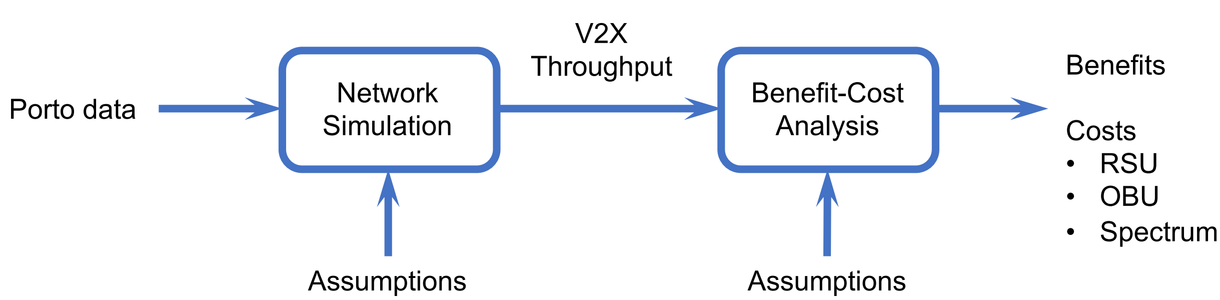

Figure 2.2. Summary of steps, inputs and outputs of the methodology.

The second step in the methodology is to use the simulated throughput to estimate benefits and costs of Internet access using V2X-based vehicular networks. Today, nearly all mobile traffic must be carried over a macrocell tower. In a capacity-limited cellular network, a reduction of traffic from mobile devices in the busy hour allows each cell to provide adequate capacity over a larger area, thereby reducing the number of towers needed to cover a given region. We define the benefit of Internet access through vehicular networks as the cost savings accrued from the difference between the number of macrocellular towers that would be necessary if there is no Internet access through V2X, and the (lower) number of towers necessary when part of the total traffic is offloaded. Benefit is compared to the costs of RSUs, spectrum and OBUs, under a wide range of numerical values for factors such as population density, penetration of OBUs in vehicles, data rate per OBU-equipped vehicle, and unit costs. Those definitions of costs and benefit are aimed to quantify how V2X-based networks for Internet access enhance social welfare, or the level of well-being in a society as a whole, independently of who incurs the costs and who derives the benefits[5]. This is described in more detail in Section 2.3.

While the quantity of onboard devices and amount of spectrum are among the definitions that characterize a scenario of analysis, the amount of infrastructure deployed for each scenario is estimated at the optimal quantity of RSUs that maximizes net benefit of Internet access over V2X-based networks. (The net benefit is defined as benefit minus infrastructure cost).

Location characteristics, i.e. whether an area is urban, suburban or rural, influence both steps. Data rates are influenced by signal propagation, which differs between urban and rural areas. Moreover, those rates are only relevant where the cellular networks are capacity-limited, which also is a condition typical for urban areas. On the other hand, those same data rates are

[5] In Chapter 4, we additionally derive cost savings for other parties, such as local governments.

expected to be higher in urban areas, because of the higher population densities. Therefore, results are likely to be more substantial in urban areas, which make them the primary focus of this analysis.

The data collected from the vehicular network in Porto, simulation model, benefit and cost model, and the numerical assumptions for the base-case scenario are described below.

2.1. Porto vehicular network and dataset

Porto is the second largest city in Portugal (Instituto Nacional de Estatistica 2011). In September 2014, its urban bus authority started offering free Wi-Fi service for passengers in 400+ buses that have OBUs equipped with an onboard Wi-Fi hotspot, as illustrated in Figure 2.3. Each bus OBU also has a router that relays traffic to/from the Internet through one of two possible paths. The preferred is through the use of DSRC, for which there were 27 RSUs (as of March 2015) deployed at fixed locations of the city. Buses were able to connect to RSUs either directly or through multihop connections using other buses. If no DSRC path is available, then data is transferred over LTE cellular networks.

Figure 2.3. Deployment of OBUs in bus in Porto, Portugal.

As of March 2015, over 2.7 TB/month were transferred. The observed volume transferred through DSRC V2X connections varies with location, with the majority of the RSUs being concentrated in downtown, where the offload ratio of bytes transferred through DSRC to the total number of bytes can reach as much as 70% at peak hours.

Taxis are also equipped with devices that collect data. Of the city's estimated total of 800 taxis, GPS positions of 400+ vehicles were collected during one month in 2012.

The data from Porto buses and taxis that were used in this dissertation are summarized in Table 2.1.

TABLE 2.1. PORTO DATA USED FOR THE ANALYSIS

Data from busses collected from October 2014 to March 2015 and from taxis in March 2012		
Data Item	Number of Observations	Description
Data volume/ position/signal per 15-second per bus	400+ buses: 240×10^6 data points	Per 15-second interval, per bus GPS position, received signal strength from RSU (if V2I-connected) or peer bus (if V2V-connected)
RSU positions	27 RSUs	Per RSU: GPS position and height
Position per second per taxi	400+ taxis: 120×10^6 data points	Per second, per taxi: time, GPS position, and an identifier of the vehicle

2.2. Network simulation

The baseline part of the simulation model represents a DSRC network of RSUs connected to the Internet, and vehicles equipped with OBUs that exchange Internet traffic with those RSUs. Transfers of packets are simulated between pairs of vehicles and between vehicles and RSUs. 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 (up to three) with other vehicles acting as relays. These flows are simulated at the transport, network, link and physical layers using the ns-3 network simulator (“Ns-3 Network Simulator” 2018).

While research such as (Lu et al. 2013) employed analytical methods to derive throughput to vehicular users, we have opted for packet-level simulation as it lets us use data from the real network in Porto to represent vehicle densities more realistically. As for the simulation platform, we have chosen ns-3, which is an open source simulation tool that has been used extensively for research on wireless networks (“Ns-3 Network Simulator” 2018; Lacage and Henderson 2006). Moreover, simulation lets us observe the impact of varying conditions such as data rates and densities of vehicles and RSUs. It would be impossible or impractical to vary those conditions in a real network such as that of Porto.

The main steps of the simulation are represented in Figure 2.4, and the main assumptions of the simulation model are described below in the following order. First, mobility and network topology, the use of ITS spectrum, the estimation of throughput rates, and endpoints for traffic flows are described. Then, the description is separated by communication layer, beginning with the transport layer and then proceeds one by one until the physical layer.

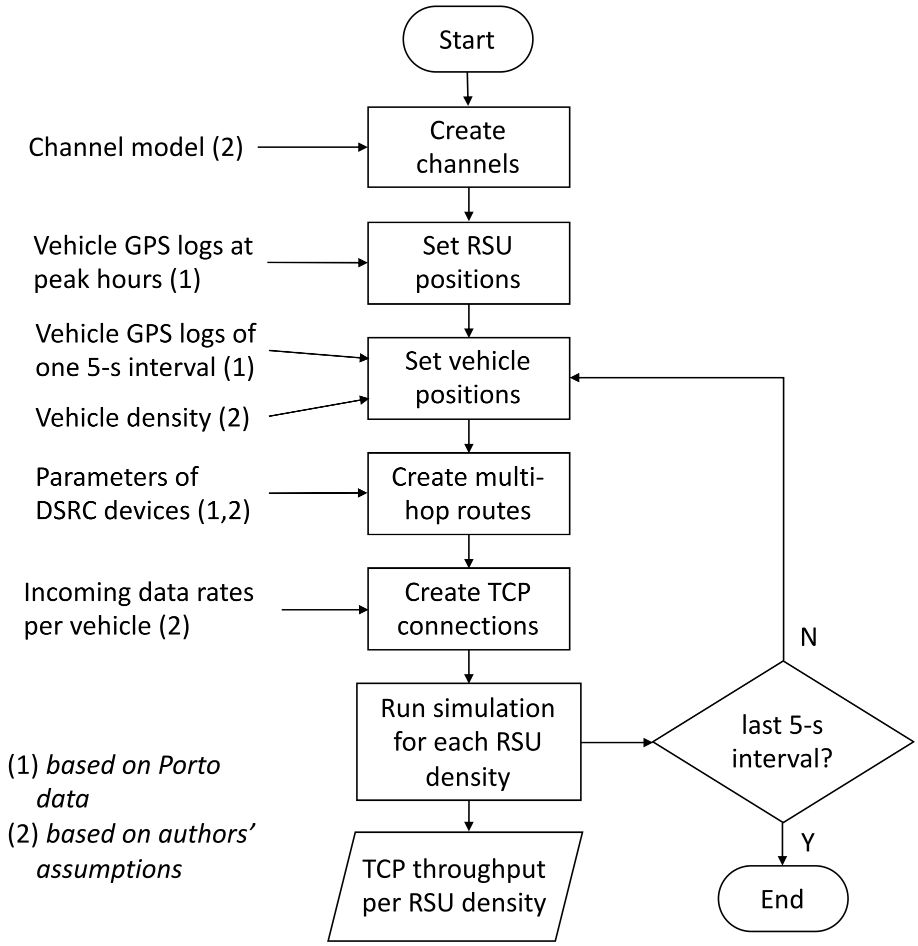

Figure 2.4. Simulation steps of V2X-based throughput for one scenario of numerical assumptions. For each scenario, throughput is simulated several times, once for each RSU density and for each 5-second interval of vehicle positions.

Vehicle mobility and RSU locations. Passenger cars, taxis and buses are positioned according to the GPS logs of taxis and buses over 20 km^2 in Porto. 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. Each of the remaining vehicles follows the same trajectory as a real taxi. This is reasonable because differently from buses, taxis do not move according to predefined routes. That is, taxis move between origins and destinations determined in an *ad hoc* basis and select routes according to congestion and driver preferences. Those taxi patterns probably apply for other passenger vehicles as well[6]. For these, we similarly select a random taxi, and a random start day and time. Vehicle and RSU antennas are placed in a tri-dimensional space. Vehicle mobility is simulated as follows. Representations of vehicles are created with static positions. Then, throughput rates between vehicles and RSUs are simulated representing a network with non-moving nodes communicating for 5 seconds. After the simulation run completes for one 5-second time interval, the process repeats for the next 5-second interval: the positions of the vehicles are changed, the simulation is run again for the referred interval, and so on. Moreover, the simulation accepts RSU density as an input variable, and then places RSUs in locations with a large number of vehicles at peak hours, using the *k-means* clustering heuristic (Moore 2001), with peak-hour vehicle positions as the input. The *k-means* algorithm divides vehicle locations among k RSUs and finds the optimal centroid for each RSU which minimizes the sum of squared distances between vehicle positions and the RSU. For each simulation scenario, the network is simulated multiple times with infrastructure density ranging from 0 to 10 RSUs/km^2.

Vehicle and RSU antennas are placed in a tri-dimensional space. Vehicle X and Y coordinates are given by the GPS data. Z coordinates represent the height of antennas. All RSU antennas have a height of 7 meters, which is the average height of Porto RSUs as of March 2015. Bus antennas have a height of 3 meters (average of buses in Porto), and all other vehicles have height of 1.5 meters (Boban et al. 2011).

[6] However, taxi patterns may not be representative of certain vehicle types such as trucks, but the total number of miles per year from passenger cars outnumber the total from trucks.

Use of ITS spectrum for Internet access. 75 MHz of DSRC spectrum is used in seven 10 MHz channels, of which one is reserved for control and management of all channels, and two others are reserved for safety applications (IEEE 2010a). We assume the four remaining channels are available for Internet access, and each vehicle OBU and each RSU is equipped with four radios. It is assumed that each packet streams flow uses as many 10 MHz channels as needed to carry the offered number of packets for that flow. The channel or channels to be used at each hop of the flow are chosen as the least used channels in the area simulated. These spectrum-related assumptions are used in Chapter 3 and Chapter 4 and are varied in Chapter 5.

Estimation of throughput of the vehicular network. The throughput rate via DSRC for each vehicle is defined as the rate achievable when that vehicle receives data from an RSU it is connected to (either through a single or multiple hops). We assume that the traffic sent downstream to a car equals the sum of the throughput over the DSRC network to that car and the throughput over the cellular network to that car. The same is assumed for traffic traveling upstream from each car. These assumptions are accurate if the amount of traffic that is lost and the amount that is unnecessarily sent on both networks are both negligible. This is reasonable as long as the cellular network is always available and has enough capacity to carry all traffic that cannot be carried over the vehicular network.

Steady-state throughput through DSRC are estimated for each 5-second interval based on the positions of all vehicles at the beginning of the interval. This simplifying assumption ignores that vehicles move during the interval, so throughput would change gradually rather than jump every 5 seconds. This form of analysis may miss some of the fluctuations in data rate as observed by a moving vehicle, but it allows for a good approximation of throughput when averaged over many time intervals as long as vehicles can switch off between the vehicular network and a ubiquitous cellular network as needed so that data rate fluctuations have little effect on the total amount of traffic sent and received. This is a reasonable first-order estimate if the time to establish links is negligible, and this switching time with DSRC is expected to be roughly 300 milliseconds (IEEE 2010a; Mussabbir and Yao 2007). To estimate steady-state throughput in a given time interval, the simulation is first run for an extended warm-up period before statistics are gathered. The warm-up part of the simulation runs for 8 seconds, and after that statistics are calculated for the data received in one second. This 8-second warm-up period was obtained by experimentation – all scenarios simulated resulted in throughput close to the mean after that period, and most do so less than 1 second after the beginning of the interval.

Each OBU-equipped vehicle is the endpoint of one and only one bidirectional flow, while each RSU may be the endpoint for zero, one or more flows, up to the number of vehicles connected to it. However, any OBU-equipped vehicle can also serve as a relay for data of a flow that has another vehicle as a destination, in case of multihop communications. Protocol-specific data include acknowledgments and retransmissions in all layers. However, those protocol messages are not considered in the statistics – only the number of application-layer data bytes received and sent by the vehicle per unit of time is considered in the throughput.

The throughput per unit of area (in bps/km^2) is defined as

$$R = \sum_{i=1}^{V} R_i \qquad (2.1)$$

where V is the density of OBU-equipped vehicles per km^2, and R_i *is* the data throughput achievable between vehicle *i* and an RSU it is communicating with.

Endpoints for traffic. Each RSU is a gateway to the Internet which a given vehicle connects to. We assume that any RSU in the network can connect to any OBU-equipped vehicle. Therefore, RSUs can offload traffic of any given vehicle, regardless the cellular operator which the user of that vehicles is subscribed to[7]. We simulate the traffic on the vehicular network, i.e., between vehicles and RSUs, so we treat the RSU as if it were the endpoint of a transport-layer connection rather than merely a gateway, as represented in Figure 2.1 (b). When addressing spectrum sharing issues we also simulate traffic originated and relayed by unlicensed devices, as described in In Chapter 5.

Transport layer. At each interval, a Transmission Control Protocol (TCP) connection is simulated between each vehicle and RSU it connects to. The TCP Maximum Segment Size (MSS) used is 2244 bytes, which is the maximum size of the packet that the 802.11 link layer supports without fragmentation (2304 bytes), minus 60 bytes for the link and IP headers (Z. Wang and Hassan 2008). That MSS is roughly similar to typical values for TCP connections traversing 802.11 networks.

Network layer. IP packets are routed through the path with the minimum number of hops between the vehicle and an RSU, up to a maximum of three hops for each path[8]. 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

[7] This assumption abstracts our model from specific conditions such as whether there are multiple cellular operators providing Internet service in a location. For example, a single Internet Service Provider (ISP) may deploy and operate RSUs in a location and have some business arrangement to offload traffic from all cellular operators. Or multiple ISPs may deploy their own infrastructure and have some business arrangement to have all RSUs connect to any vehicle.

[8] Multihop communications may result in significant mutual interference between adjacent hops in a path. When OBUs are equipped with multiple radios and have access to multiple channels, we allocate channels to paths such as adjacent hops use distinct channels. However, we have found that the overall throughput *R* with this configuration is not much different from the throughput *R* obtained with all hops in a path using a single channel, especially for the base-case data rate and vehicle density. One reason is that most vehicles in this scenario that are connected do so via single-hop paths.

one path randomly among the paths with the minimum number of hops, such that each of those paths is equally likely.

Link layer. The media access control (MAC) sublayer in the DSRC link layer is the one specified in the IEEE 802.11p amendment (IEEE 2010b) of the IEEE 802.11 standards.

Physical layer. A hop is used between two nodes only if signal strength at the receiver exceeds 15 dB above the sensitivity threshold (-94 dBm). This is the criteria determined empirically in the vehicular network of Porto. When the hop is used, the transmitter selects the modulation scheme (among the schemes defined in IEEE 802.11p) in accordance to the signal strength at the receiver. (However, only the highest modulation schemes are chosen, because of 15 dB criteria mentioned above.) Packets are received at an error rate that also depends on the signal-to-interference-plus-noise ratio (SINR), as described in (Lacage and Henderson 2006) and (“Ns-3 Network Simulator” 2018).

The transmitted power is 14.6 dBm, obtained from measurements at the equipment output, which is also consistent with (Cardote et al. 2012) and (Bai, Stancil, and Krishnan 2010), and the gains of the transmission antennas are 16 dBi and 5 dBi for the RSUs and vehicles, respectively, which are consistent with the settings of the equipment in the Porto network.

The received signal is calculated according to the propagation loss model from (Meinilä et al. 2009). (urban microcell B1 variant). This model is appropriate for urban areas, because vehicular networks are most useful in urban areas where density of vehicles is higher, as is demand for cellular networks. Moreover, that model was preferred over other similar models because it is valid for the ITS band (5.9 GHz), and it explicitly models two other characteristics that are relevant in vehicular networks: whether those nodes are in line-of-sight (LOS) or non-

LOS (NLOS) (Meinilä et al. 2009; Zhao et al. 2006), and the antenna heights of vehicles and RSUs (Mecklenbrauker et al. 2011; Meinilä et al. 2009). For LOS, the loss *L* is given (in dB) as

$$L = PL_{LOS} + N{\sim}(0, \sigma) \qquad (2.2)$$

where

$$PL_{LOS} = L_0 + 10\, n\, log_{10}(d) \qquad (2.3)$$

is the path loss calculated as a reference loss L_0 and a function of the distance *d* (meters), and the path loss exponent *n* represents the degree of attenuation. *N* is a Gaussian random variable with zero mean and represents large-scale fading effects such as shadowing of the LOS path by obstacles. For LOS the values are

$$n = \begin{cases} 2.27\ for\ d < d_{BP} \\ 4 \quad\ for\ d \geq d_{BP} \end{cases} \qquad (2.4)$$

$$L_0 = \begin{cases} 41 + 20\, log_{10}\left(\dfrac{f}{5 * 10^9}\right) & for\ d < d_{BP} \\ 9.45 - 17.3\, log_{10}(h_1 - 1) - 17.3\ \ log_{10}(h_2 - 1) + 2.7\, log_{10}\left(\dfrac{f}{5 * 10^9}\right) & for\ d \geq d_{BP} \end{cases}$$

$$\sigma = 3 \qquad (2.5)$$

where *f* is the ITS frequency in Hz, h_1 and h_2 are the heights of the vehicles and/or RSU, and

$$d_{BP} = 4(h_1 - 1)(h_2 - 1)f/c \qquad (2.6)$$

where *c* is the speed of light in m/s.

For NLOS,

$$L = PL_{LOS} + 20 - 12.5\, n + 10\, n\, log_{10}(d) + 3\, log_{10}\left(\frac{f}{5 * 10^9}\right) + N{\sim}(0, \sigma)$$

$$n = max(2.8 - 0.0024\, d, 1.84)$$

$$\sigma = 4 \qquad (2.7)$$

Other models (Zhao et al. 2006) provide similar path loss and shadowing parameters not substantially different from those shown above.

Each interval each link is assumed to be in LOS or NLOS according with probability *Prob(LOS)* estimated as (Calcev et al. 2007)

$$Prob(LOS) = \begin{cases} \frac{d-300}{300} & for\ d < 300 \\ 0 & otherwise \end{cases} \qquad (2.8)$$

(Asplund et al. 2006) and (Meinilä et al. 2009) propose expressions which results similar LOS probability.

In addition to path loss and shadowing, some models include zero-mean random variables to represent fast-fading effects such as multipath propagation and Doppler spread (Mecklenbrauker et al. 2011). In our simulation model, the estimated path loss and shadowing components are assumed to be constant over each 5-second interval, and the effect of fast-fading is assumed as negligible, as we estimate average losses across many links rather than predict the loss of a particular link.

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 is a reasonable approximation for the observed loss. For example, at a distance of 100 m between an 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.

2.3. Benefit-cost analysis

The second step of the methodology is to use DSRC throughput to estimate benefits and costs of Internet access at peak-hours. Our benefit-cost analysis quantifies the net economic impact of deploying V2X-based networks to offload traffic from existing cell towers. Benefit-cost analysis is a widely recognized method that has been employed for decades by the U.S. federal government and many other entities (U.S. Department of Transportation 2016; Boardman et al. 2001) to assess the social impact of new policies.

Our basic definitions of costs and benefits are independent of who incurs those costs and who derives those benefits. This allows us to quantify the impact of deploying DSRC infrastructure on total social welfare[9] without making any assumptions about things like who pays for RSUs, how much the operator of RSUs charges for the service, and who pays for the service. In Chapter 4, we additionally consider the implications of costs being incurred by Internet Service Providers and government agencies responsible to deploy infrastructure.

We define the benefit of Internet access through vehicular networks as the net present value (NPV) of cost savings, under the following assumptions. All macrocellular carriers in the region being analyzed are assumed to be capacity-limited, which means the system is expected to operate at full capacity at peak hours. Therefore, Internet usage in vehicles as new mobile traffic should be met either via capacity expansion of the macrocellular networks, or via offload as represented in Figure 2.5. To serve more users or higher rate per user, a capacity-constrained carrier that is already using current technology throughout the spectrum available to it must deploy new towers. Besides tower deployment, there are two other ways to increase macrocellular capacity: addition of spectrum, or changing the efficiency of the technology, such as replacing older equipment with newer systems, or adding sectors per tower. Since network designers will generally choose the approach that is most cost-effective at the time, the marginal cost of increasing capacity is likely to be similar for all available approaches (Tan and Peha

[9] Social welfare is the level of well-being in a society. In economic theory, it is the sum of all benefits experienced, minus the sum of all costs incurred to provide those benefits, regardless of who benefits and who incurs the costs.

2015). We assume that the deployment of new towers is the preferred method to increase macrocellular capacity. Carriers do deploy towers when they need capacity, in part because spectrum is costly to obtain and typically takes a long time from allocation to commission. Moreover, carriers that need more capacity in a region are often already using current technology there (Clarke 2014) and have often deployed the maximum possible number of sectors.

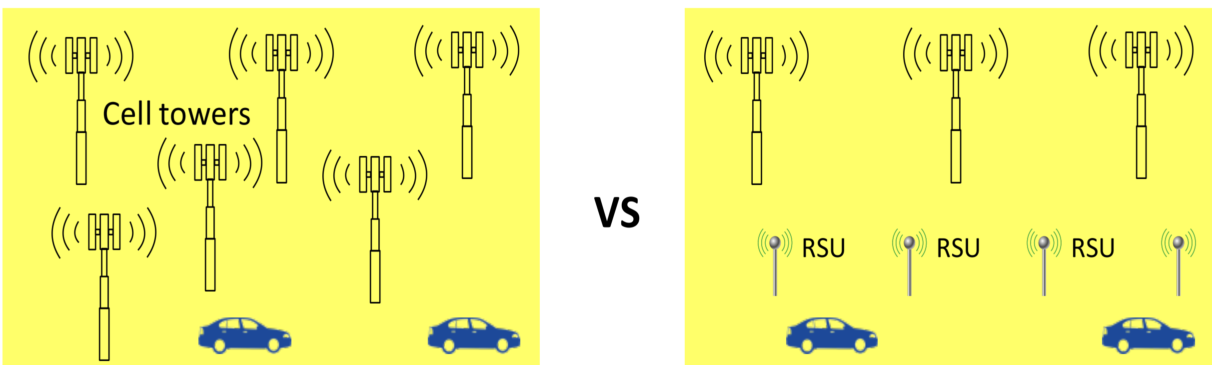

Figure 2.5. The economic benefit is the difference between the cost to provide Internet access for vehicles using only macrocellular towers, and the cost to carry part of the Internet traffic through the V2X-based network.

It is assumed that devices will send as much traffic as possible over the V2X network. The amount of traffic carried through V2X equals the reduction in the amount of traffic carried through cellular, i.e. every bit carried on the vehicular network is one less bit on the cellular network. This means that devices switch between the V2X and macrocellular networks with negligible disruption, with no data being lost or transmitted in duplicity through both networks.

The NPV of the benefit of Internet access per km^2 is

$$NPVB = \rho_{savedtowers} * C_{tower} \qquad (2.9)$$

where $\rho_{savedtowers}$ is the total number of macrocell towers "saved" per unit of area due to Internet access through vehicular networks and C_{tower} is the average NPV per macrocell tower. We assume that if there is sufficient capacity downstream then there is also sufficient capacity upstream, and that carriers are using Frequency Division Duplexing (FDD) so spectrum can be labeled as either upstream or downstream. In a cellular network, the maximum downstream capacity $bps\_max_{km^2}$ in bits per second per unit of area is given by

$$bps\_max_{km^2} = s_{sector}\frac{bw}{FR}\rho_{towers}\,N_{sectors} \qquad \textit{(2.10)}$$

where s_{sector} is the average downstream spectral efficiency in bits per second per hertz per sector[10], *bw* is the bandwidth per macrocellular carrier used for downstream transmission, $FR \geq 1$ is the frequency reuse factor, ρ_{towers} is the number of towers per km^2 and $N_{sectors}$ is the number of sectors per tower.

In order to serve all fluctuations of demand, the maximum capacity should equal or exceed the data rate demand at peak hours. Therefore, if $bps1_{km^2}$ is the peak-hour, downstream rate demand per unit of area from macrocells when no Internet access through V2X-based networks takes place, the number of towers necessary $\rho 1$ per unit of area is found by solving the equation below for $\rho 1$:

$$bps1_{km^2} = s_{sector}\frac{bw}{FR}\rho 1\,N_{sectors} \qquad \textit{(2.11)}$$

[10] The efficiency in bits per second per hertz of a given cellular user depends on the signal-to-interference-and-noise ratio experienced by that user. In general, users located at the edge of a cell will experience less efficiency than users located closer to the cellular tower. Nevertheless, in our model we take an average estimate of the spectrum efficiency across all users in a sector, since we are interested in the overall capacity of each cell rather than particular users.

Let $bps2_{km^2}$ also be the data rate demand from macrocells, but when part of the traffic is carried through V2X networks, to be served by $\rho 2$ towers. The difference between $bps1_{km^2}$ and $bps2_{km^2}$ is the traffic offloaded per unit of area:

$$bps1_{km^2} - bps2_{km^2} = bpsOff = s_{sector} \frac{bw}{FR} \rho_{savedtowers} N_{sectors} \qquad (2.12)$$

then the number of saved macrocell towers is

$$\rho_{savedtowers} = \frac{bpsOff\ FR}{s_{sector}\ bw\ N_{sectors}} \qquad (2.13)$$

where *bpsOff* is the downstream throughput over the V2X network.

The total cost of Internet via V2X-based networks per km^2 $NPVC$ consists of three types of costs:

$$NPVC = NPVC_{RSU} + NPVC_{OBU} + NPVC_{Spectrum} \qquad (2.14)$$

where $NPVC_{RSU}, NPVC_{OBU}$ and $NPVC_{Spectrum}$ are the NPV per km^2 of the costs of RSUs, OBUs and ITS spectrum, respectively, and are given as

$$NPVC_{RSU} = \rho_{RSU}\ C_{RSU}$$

$$NPVC_{OBU} = \rho_{OBU}\ C_{OBU}$$

$$NPVC_{Spectrum} = \rho_{Spectrum}\ C_{Spectrum} \qquad (2.15)$$

where ρ_{RSU} is the number of RSUs for Internet access deployed per unit of area, which is assumed to be independent and not shared with RSUs deployed for safety or purposes other than Internet access, ρ_{OBU} is the number of OBUs deployed per unit of area, $\rho_{Spectrum}$ is the amount of DSRC spectrum in MHz times the population density, and $C_{RSU}, C_{OBU}, C_{Spectrum}$ are the NPV per RSU, OBU, and MHz of spectrum per person (also known as the cost per MHz-pop), respectively.

The comparison between the benefit and costs defined above depends on the decision to be made, i.e. some of the costs that are relevant for one decision may be irrelevant for another. For example, when evaluating the cost-effectiveness of deploying V2I infrastructure in the context of a safety mandate, ITS spectrum is allocated and OBUs are purchased for safety reasons[11]. In this case, spectrum and OBU costs are sunk with respect to Internet access, and a decision to deploy RSUs increases social welfare if and only if benefit of Internet via V2X-based networks exceeds RSU costs. However, if there is no safety benefit derived from the mandate, then ITS spectrum and OBU costs are not sunk, and social welfare will increase only if benefit of Internet access exceeds all V2X-related costs: RSUs, OBUs and ITS spectrum.

2.4. Base-case scenario

The base-case numerical values are listed below for the assumptions used in the estimates of the throughput of Internet access via V2X-based networks, and the associated benefits and costs. Those assumptions apply for the results in Chapters 3, 4 and 5, unless otherwise stated. In those chapters, some of those numerical assumptions are varied either to address specific research questions or for sensitivity analysis.

Monetary values. The monetary values are in constant 2014 dollars. Benefit and cost NPVs are calculated at a real discount rate of 7% over a horizon of 10 years. The discount rate

[11] DSRC technology enables an OBU to send safety and non-safety messages simultaneously with distinct priority (Kenney 2011). Therefore, it is reasonable to assume that OBUs deployed for safety applications can be used also for non-safety communications with negligible additional cost.

is consistent with the rate recommended by the U.S. Office of Management and Budget for benefit-cost analysis of federal programs (Office of Management and Budget 1992). Other analyses use similar rates (Hallahan and Peha 2009; Harding et al. 2014). The 10-year horizon is long enough to evaluate the lifetime costs of the main elements of the model. For example, RSU lifetime was estimated to be 10 years in analysis for the U.S. Dept. of Transportation (Wright et al. 2014). Although some costs such as macrocellular towers are incurred for a longer horizon, because of the 7% discount rate, their NPV is primarily determined in the first 10 years (especially if the upfront part of the costs is large when compared to the recurring part).

Population density. We make the simplifying assumption that population density is constant throughout the region being analyzed.

Number of vehicles on the road at peak hours per capita. Assumed as in Table 2.2, which is calculated as the product of vehicles owned per capita (United States Census Bureau 2015), fraction of time vehicle is in use and ratio of peak-hour usage to average usage (Santos et al. 2011). We consider usage at peak hours because our calculation of benefit is based on data offload from capacity-limited cellular networks, and it is peak-hour usage that determines how much capacity a cellular carrier needs, and thus the cost that the carrier incurs.

TABLE 2.2. NUMBER OF VEHICLES ON THE ROAD AT PEAK HOURS PER CAPITA, AS A FUNCTION OF POPULATION DENSITY

People per km^2	Number of vehicles owned per capita	Number of vehicles on the road at peak hours per capita
10	1	0.077
200	0.8	0.061
1000	0.65	0.050
2000	0.6	0.046
3000	0.55	0.042
5000	0.44	0.034
12000	0.22	0.017

OBU Penetration in vehicles. Assumed as 25%. This is reasonable for a decision-maker looking 5 to 10 years ahead in the context of a mandate to deploy V2X devices in all new light vehicles (U.S. Department of Transportation 2016; Harding et al. 2014). (Considering that cars in the U.S. have an average life cycle in the order of 10 years, it is reasonable to expect that 8-10% of the fleet is renewed each year.)

Data traffic per DSRC-equipped vehicle on the road. At any 5-second interval during the peak hour, 50% of the OBU-equipped vehicles on the road are endpoints for data being 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. This is consistent with predictions that vehicular traffic will reach 5 GB/month per vehicle in the coming years (Deutsche Telekom 2013). In reality, data rates vary from vehicle to vehicle at any given time, but since RSUs are typically in range of dozens of OBU-equipped vehicles at all times during peak hour, this simplifying assumption should have limited effect on aggregate throughput.

Share of downstream traffic. While a vehicle is transferring data, 90% of the data flows in the downstream direction (RSU to vehicle). In the Porto V2X-based network, 92% of a session volume is downstream, on average, and (Sandvine 2014) reports a similar ratio for the monthly usage per mobile device in the U.S.

Use of DSRC spectrum for Internet access. The band allocated for DSRC in the U.S. is divided in seven 10 MHz channels, of which three are reserved for safety applications and for control of operation of the other channels (Campolo and Molinaro 2013). We assume the four remaining channels are available for non-safety traffic, and each OBU and RSU is equipped with four radios. Multiple radios allow different channels to be used simultaneously, so multi-radio OBUs can increase channel utilization and throughput compared to single-radio OBUs. Therefore, if RSUs that act as Internet gateways are deployed, and cars must be equipped with OBUs anyway, then users have much to gain by getting multi-radio OBUs rather than single-radio OBUs. Moreover, we believe the cost difference is likely to be quite small. For example, Chen et al. suggest that if certain technical problems can be solved, then this cost will be small, and "there is every reason" to expect the use of more radios than is strictly necessary (Chen, Jiang, and Delgrossi 2009). Nevertheless, it is possible that OBUs with fewer radios will be deployed, which would lead to somewhat lower throughputs than we report in this paper.

Unit cost of macrocellular tower. The NPV of cost per macrocell tower over 10 years is $750,000. Where carriers are leasing space on existing cell towers, this cost includes leasing fees. Where carriers build their own towers, a decade of leasing fees is replaced by upfront capital expenditures (CAPEX). A 10-year NPV of $750,000 is roughly consistent with previous estimates (U.S. Federal Communications Commission 2010; Hallahan and Peha 2011a; Newman 2008), in 2014 dollars.

Macrocellular spectrum efficiency. The downstream average efficiency of a macrocell is 1.4 bps/Hz/sector, which is an accepted value for LTE-FDD rel. 8 (Sesia, Toufik, and Baker 2011). Some devices will be more spectrally efficient, such as those using newer 3GPP cellular standards, while usage of less efficient devices also continues, sometimes with efficiencies below 1 bps/Hz/sector (Clarke 2014).

Sectors per macrocell. Each macrocell is divided in 3 sectors, which is consistent with (Sheikh 2014) and others.

Macrocellular bandwidth. Any new tower deployed in a capacity-limited region is constrained by the bandwidth available for downlinks, and we assume that each tower would operate over a downlink bandwidth of 70 MHz per sector. A tier-1 provider is estimated to hold roughly 30 MHz of downlink spectrum for LTE, on average (Goldstein 2015), and spectrum in use for LTE is estimated at about half of total spectrum for mobile broadband. Substantial amounts of new spectrum are expected to be allocated for newer cellular technologies, but its effective use may take several years for actual deployment.

Macrocellular frequency reuse factor. The frequency reuse factor in macrocells is 1, which is consistent with a typical macrocellular network configuration with current technology (Wannstrom and Mallinson 2014).

Unit cost of RSU. The average NPV over 10 years of an RSU deployed for V2I communications is $14,000. This is based on U.S. DOT estimates of average annual cost between $2,000-3,000 (Wright et al. 2014), which include replacement costs every 5 to 10 years. However, in Chapter 3 we will consider variations of more than 50% from the base-case value, as conditions about infrastructure availability may vary. For example, the city of Porto deployed RSUs for a Capex of between $1,200-4,000, by placing RSUs in existing structures (traffic poles, buildings, etc.) already owned by the city and already equipped with energy and backhaul access. The cost per RSU could be also be lower if RSUs deployed for Internet access are shared for safety or vice-versa, although sharing depends on many issues. These issues include whether the optimal placement of RSUs for Internet access matches the placement for safety communications and whether devices for Internet access and devices that are safety-critical are placed under shared control. These issues are addressed in Chapter 4. For the base-case value of the cost per RSU, no sharing is assumed. On the other hand, costs can be significantly higher if new poles, energy and communications infrastructure must be built entirely for RSUs deployed exclusively for Internet access.

Chapter 3. Throughput and Economics of Internet of Vehicles

3.1. Introduction

Internet access for users moving at high speeds, such as those in vehicles, have typically been provided through cellular networks. It may be expensive to expand the capacity of cellular networks to meet the sharp growth that is expected in mobile Internet traffic, especially in urban areas. If part of that traffic is offloaded onto other networks at a lower cost than expanding cellular infrastructure, then social welfare is enhanced by providing cheaper Internet access. In this chapter we show that V2X-based, vehicular mesh networks could be an important new way to provide Internet access in urban areas. We examine throughput, benefits and costs of Internet access via V2X-based networks for a location of homogeneous population density in the context of a mandate of OBU devices in vehicles, with a given amount of ITS spectrum being used to carry Internet traffic, and with an ISP deploying RSUs exclusively for Internet access. Under those conditions, we determine the range of population densities over which V2X-based networks are cost-effective for Internet access when compared to in-vehicle Internet traffic being carried over microcellular networks. We also show how benefits and costs depend on factors that vary over time, such as the penetration of OBU devices and data volume of Internet traffic, and uncertain factors such as unit costs of V2X and macrocellular equipment, and bandwidth.

In some circumstances, DSRC technology is more cost-effective than expanding the capacity of cellular networks. We determine what those circumstances are using the engineering-economic approach described in Chapter 2, in which we combine our simulation model with data collected from the actual vehicular network that is operating in Portugal. We use the model to estimate how much Internet traffic can be offloaded to vehicular networks that would otherwise be carried by cellular networks, under a variety of conditions. These conditions include the density of vehicles and RSUs, and data rates of incoming traffic. Moreover, we examine how throughput scales as those factors increase over time.

We use simulated throughput to estimate the benefits of cost savings of reduced cellular infrastructure due to offload, and the cost of the V2X-based network to carry that traffic. Then, we determine when benefit exceeds cost. Such a cost-benefit analysis informs important decisions regarding whether resources should be invested in vehicular networks for Internet access, rather than just vehicular safety. One decision is whether to invest in RSUs for Internet access. We find that deployment of RSUs in dense urban areas is likely to increase social welfare fairly soon after a mandate to put OBUs in vehicles becomes effective. Moreover, we find that deployment will increase social welfare in less densely populated areas over time, as the penetration of DSRC in vehicles and data rates increase.

Other decisions include whether to allocate ITS spectrum and mandate OBUs in the first place, if these steps are not taken for safety reasons. In situations in which benefit of Internet access exceeds all types of DSRC cost, then social welfare is increased by mandating DSRC devices in all vehicles and allocating spectrum regardless of whether there are benefits other than Internet access. We find that the benefits from Internet traffic alone are not enough to justify a universal mandate to deploy DSRC in all vehicles, i.e., the benefits of Internet access alone are less than total costs including RSUs, OBUs and ITS spectrum. However, the majority of V2X-related costs must be incurred anyway if safety is to be enhanced, and the associated benefits from lives saved, mitigated injures and reduces property damage are to be experienced.

Part of the findings discussed in this chapter is published in the IEEE Access journal (Ligo et al. 2017). Moreover, parts were presented at TPRC – Research Conference on Communications, Information and Internet Policy (Ligo et al. 2015), at the IEEE Vehicular Technology Conference (Ligo, Peha, and Barros 2016), and at the MASITE-ITSPA (Mid-Atlantic Section of The Institute of Transportation Engineers – Intelligent Transportation Society Pennsylvania) Conference (Ligo and Peha 2016).

This chapter is organized as follows. Section 3.2 describes previous research and how our work is positioned. Section 3.3 describes how the engineering-economic method described in Chapter 2 is employed for the analysis in this chapter. Section 3.4 contains the results, and Section 3.5 ends the chapter with the conclusions.

3.2. Related work

Previous work discussed issues related to vehicular communications over heterogeneous networks. In the survey in (Hossain et al. 2010), the authors predicted that ubiquitous deployment of DSRC may take decades, and therefore OBUs that switch between DSRC and cellular are possibly cost-effective solutions. In this chapter, we examine the actual conditions under which DSRC is cost-effective when compared to cellular, by quantifying the economic benefits and costs of offloading cellular traffic through DSRC. Other work focused on specific technical issues such as the method to select among heterogeneous networks (Tian et al. 2015; Zhang, Sirbu, and Peha 2017). In contrast, we determine the locations where it is cost-effective to offload traffic from cellular and would constitute a choice between heterogeneous networks. That is, we assume that DSRC networks for Internet access will be deployed only where they are cost-effective, and then we quantify benefits and costs at those locations. (We also assume that QoS will be satisfied by whatever network that is selected.)

There is extensive work on the technical capabilities of DSRC-based vehicular communications, e.g. (Kenney 2011; Mecklenbrauker et al. 2011; Uzcategui and Acosta-Marum 2009; Campolo and Molinaro 2013; Zeadally et al. 2010), and economic benefits and costs of DSRC technology for safety communications (U.S. Department of Transportation 2016). However, to our knowledge there has been little work on the cost-effectiveness of vehicular communications over heterogeneous networks for non-safety applications. The leading exception (Lu et al. 2013) compared costs of various architectures when deploying greenfield infrastructure that would provide ubiquitous coverage in a given region using a given 10-MHz block of spectrum. This is related but different from the scenario we address, in that we assume that cellular carriers already provide ubiquitous coverage in cellular spectrum, and the question is whether it is more cost-effective to expand existing cellular carriers or deploy infrastructure for vehicular networks operating in the ITS band.

In (Lu et al. 2013), the authors compare three types of infrastructure: cellular towers that provide macrocells, “roadside access points” that provide microcells, and mesh networks. For each type, direct communications between mobile devices and infrastructure is supplemented with V2V vehicular communications if and only if the infrastructure density is insufficient to provide ubiquitous coverage. In these cases, the authors assume that there will be enough vehicles to cover all gaps in coverage. For each type of infrastructure, lower bounds of throughput capacity are derived as a function of infrastructure density, and costs are compared for a fixed capacity. When the desired capacity is low, they conclude that roadside microcells are less cost-effective than macrocells and mesh infrastructure. However, if the desired capacity is higher they conclude that roadside microcells are more cost-effective than macrocells. This is a somewhat surprising result from (Lu et al. 2013), considering that current greenfield deployments for mobile Internet service typically start with macrocells rather than microcell or mesh infrastructure.

The findings in (Lu et al. 2013) are in part the result of assumptions that are somewhat different from those of this chapter. For example, the authors assume that macrocellular networks have a frequency reuse factor of 9, and no cell sectorization. In contrast, we assume a reuse factor of 1 and 3 sectors per cell, as we might expect in an urban LTE deployment. They assume that either cellular or vehicular networks would operate in the same 10 MHz block. In contrast, we assume that cellular carriers operate in 70 MHz of spectrum as is typical for a large provider, and vehicular networks operate in 40 MHz of spectrum at a much higher frequency in accordance with FCC regulations for the ITS band. The authors in (Lu et al. 2013) assume frequency reuse can be managed so that there are no packet collisions, even in a vehicular network which can have hidden terminals. To take the impact of collisions into account, as well as congestion, we use packet simulation with protocols and parameters consistent with the DSRC, IP and TCP standards.

Some carriers and researchers are considering the use of fixed Wi-Fi hotspots that offload vehicular data traffic that is tolerant to delays (AT&T 2015; Comcast 2013, 2015; Balasubramanian, Mahajan, and Venkataramani 2010; Eriksson, Balakrishnan, and Madden 2008; K. Lee et al. 2010; Balasubramanian et al. 2008). Moreover, there has been research on the resulting economic impact of Wi-Fi offloading (Markendahl 2011; J. Lee et al. 2014). However, vehicular networks offer new opportunities for Internet access that are quite different from what is possible with Wi-Fi hotspots, and this requires new analysis.

The benefits of vehicular networks are different from Wi-Fi hotspots because the traffic carried is different. Wi-Fi is often a good solution for users who are stationary for the period when they are accessing the Internet, but it is often inadequate for users who access the Internet while moving. In addition, the costs associated with vehicular networks are quite different from the costs of typical Wi-Fi networks, which are generally microcellular. In a V2X-based mesh network as illustrated in Fig. 2.1, a relatively small number of RSUs can connect a large number of vehicles equipped with OBUs to the Internet. It also helps that DSRC links can be longer than typical connections to Wi-Fi hotspots, i.e. up 250-350 meters in cluttered urban areas, as measured in Portugal. Although far fewer fixed devices are needed to cover an area with a vehicular network than with Wi-Fi, those fixed RSUs are also more expensive, because they must operate outdoors in hostile conditions, and they are not currently mass produced. Because of these differences, this chapter is important as it examines the cost effectiveness of V2X-based mesh networks to offload Internet traffic.

3.3. Methodology

We use the engineering-economic approach based on network simulation and benefit-cost analysis, as described in Chapter 2, to evaluate the cost-effectiveness of the V2X-based networks for locations of homogeneous population density in the context of a mandate of OBU devices in vehicles, with a given amount of ITS spectrum being used to carry Internet traffic, and with an ISP deploying RSUs exclusively for Internet access. In this section, we discuss the characteristics of the methodology that are specific to this chapter.

As described in Chapter 2, the costs of OBUs, RSUs and ITS spectrum are proportional to their quantities. The quantity of OBUs is assumed as dependent on population density, but independent from the unit cost of an OBU. The NPV of the cost of an OBU is assumed as $350. This is based on U.S. NHTSA estimates (Harding et al. 2014) and considering four radio interfaces and antennas per vehicle. For the base-case scenario the population density is chosen as 5,000 people/km^2, which is representative of the city of Porto (5,600) (Instituto

Nacional de Estatistica 2011) as well as large cities such as Boston or Chicago (United States Census Bureau 2015).

On the other hand, the quantity of RSUs considered optimal in a scenario does depend on the cost per RSU. This is because the optimal RSU quantity is considered as the one that maximizes the difference between total benefit and total RSU cost. In this chapter, the cost per RSU and the derived results reflect a scenario where RSUs are deployed for Internet access only. In contrast, in Chapter 4 we evaluate how benefits, costs and other measures are affected when RSUs are shared for purposes beyond Internet access.

Moreover, in this chapter the amount of ITS spectrum used for non-safety purposes is assumed as fixed and independent from the cost of that spectrum. We this cost as $0.10 per MHz per capita (MHz-pop). This value is uncertain, as the cost of spectrum depends on frequency (Kerans et al. 2011; Tan and Peha 2015; Alotaibi, Peha, and Sirbu 2015; Peha 2013b), and the market value above 5 GHz is not well-established. The cost per MHz-pop above that is used in this chapter is a lower bound, which is revisited in Chapter 5 where a wide range of values well above and below $0.10 is discussed.

Several of the other numerical assumptions presented in Chapter 2 are varied in this chapter. The objectives are twofold. First, we discuss how the cost-effectiveness of V2X-based networks is going to evolve over time, as factors such as data rates and OBU penetration increase as some predict. Second, we evaluate how the cost-effectiveness derived from the base case scenario changes with the uncertainty on the numerical assumptions. The numerical assumptions that are varied in this chapter include population density and penetration of OBU devices in vehicles, which both determine the quantity of OBU-enabled vehicles on the road, data rates of incoming Internet traffic, which together with the quantity of vehicles influence both the interference among V2X devices and throughput, and unit costs and bandwidth of microcells, which affect our estimates of the benefit of offload.

3.4. Results

This Section presents the simulated DSRC throughput, benefit and cost results for the base case scenario, and how those results vary if base case values change. The network simulation is run with as many vehicles and for as many 5-second intervals as necessary such as the resulting statistical significance is sufficient to support conclusions. We average throughput over 10 time-intervals for at least 1,000 vehicles in a 20 km^2 region. If we make the simplifying assumption that the throughputs of these 1,000 vehicles are mutually independent, although throughputs at different time intervals are not, then mean throughput is 170 kbps, which is about 40% of the base case data rate per vehicle described in Chapter 2. The 95% confidence interval for the simulated throughput of a given scenario of device densities and data rates is within 7% of the mean.

3.4.1. Base-case scenario

Figure 3.1 shows throughput as a function of RSUs per km^2 under base case assumptions. Throughput increases with more infrastructure, as the number of vehicles that reach an RSU increases. However, the marginal gains in offload rate decrease as the RSU density exceeds 2 per km^2. This matters because while increasing RSU density increases DSRC throughput and therefore benefit, it also increases cost. This can be seen in Figure 3.2, which shows both benefits and costs as a function of RSU density under the same assumptions. Figure 3.2 shows that for the base case values, the maximum difference between benefits and costs occurs at 1 RSU/km^2, for which benefits exceed the cost of RSUs by 50%. If spectrum has already been allocated and OBUs purchased, as is likely to occur for safety applications, then those are sunk costs. Consequently, the benefits of deploying RSUs exceed the costs, and doing so will increase social welfare. However, Figure 3.2 also shows that the benefit of Internet access is considerably less than the cost of OBUs. Thus, the value of deployment of vehicular networks for Internet access alone, i.e., without consideration of the improvements in highway safety, are not sufficient to justify the deployment of OBUs and the allocation of spectrum in the base case scenario.

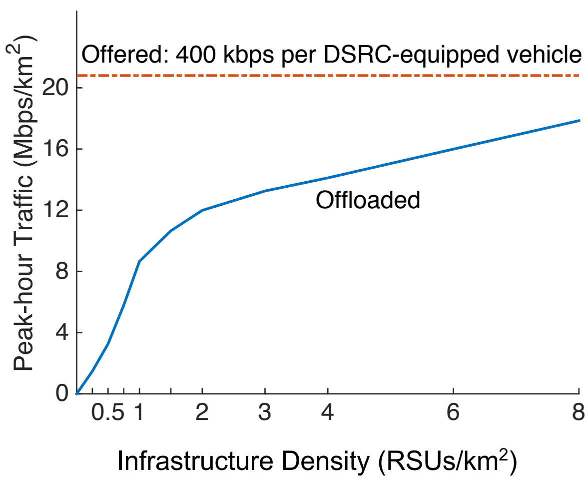

Figure 3.1. Average traffic offered and offload throughput rate at a peak hour, for the base case scenario.

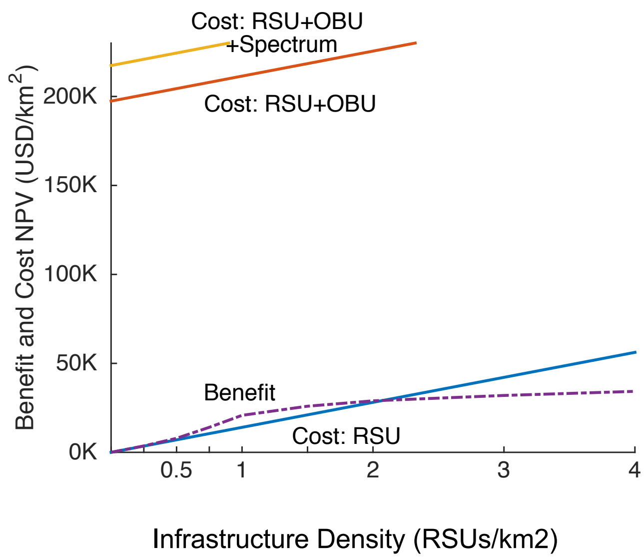

Figure 3.2. Benefit and cost for varying infrastructure density, for the base case scenario.

3.4.2. Impact of population density

The previous subsection showed that deploying RSUs can increase social welfare in the baseline case, which corresponds to a densely populated city such as Porto or Chicago. However, that may not be the case everywhere. In a more densely-populated area, there will be a greater density of vehicles and more on the road at peak hours. Therefore, more in-vehicle OBUs will be used, and more RSUs will be deployed for those vehicles to connect to, so OBU and RSUs costs increase with population density. On the other hand, throughput per unit of area, and hence benefit, are also expected to increase.

Figure 3.3 and Figure 3.4 show throughput, benefits, and costs as a function of population density. Traffic per vehicle, penetration, unit costs and spectrum parameters are held constant at base case values. The benefit and cost of RSUs in Figure 3.4 depend on the optimal quantity of RSUs for each population density, which is chosen as described in Chapter 2. For the values of population density in which the NPV of benefit of Internet access exceeds the NPV of cost of RSUs, the number of RSUs chosen is the quantity that maximizes the difference between the NPV of benefit and the NPV of RSU cost.

For the population density values in which the NPV of benefit is lower than the NPV of RSU cost for any quantity of RSUs, we calculate the quantity of RSUs as a linear extrapolation from the population density range which the NPV of benefit is greater than the NPV of RSU cost. This calculation is only to illustrate that deploying RSUs is *not* cost-effective for the population densities where benefit is less than RSU cost. For those populations the optimal RSU density is actually zero, and so are benefit and RSU cost.

Figure 3.3 shows that data rates of incoming traffic (labeled as “offered” in the graph) increases rapidly as a function of population density, which is expected because quantity of vehicles increases with population density. V2X-based throughput also increases with population density, although at a slower pace than offered traffic because competition for the use of the wireless medium limits offload.

Figure 3.4 indicates that benefit increases faster than RSU cost. The reason is that throughput grows roughly proportionally to population but the optimal number of RSUs rises at a slower pace. For base case assumptions, the threshold for which benefit exceeds cost is 4,000 people/km^2. If decisions about whether to deploy RSUs are made on a city-wide basis, this means cities with population densities at least as great as Chicago or Porto would benefit from RSU deployment, assuming there is already spectrum allocated and a mandate of OBU devices for safety purposes. However, RSUs could be deployed within an area much smaller than a city, and many cities with more modest population densities have some neighborhoods with densities over 4,000 people per km^2.

Figure 3.4 shows that benefit grows faster than RSU cost, but OBU cost grows much faster than benefit. Under a mandate, every vehicle will (eventually) incur OBU costs, but only the vehicles on the road at peak hours add to the benefits. However, if adoption is voluntary, then OBU penetration is lower and total OBU cost would be less than under a mandate. Moreover, owners of vehicles that are often in use are more likely to adopt, and this would also have the effect of increasing the ratio of OBU-equipped vehicles on the road at peak hour to total cars. Thus, if many of the OBU-equipped cars are driven extensively, then this will also increase the net benefit of deploying RSUs.

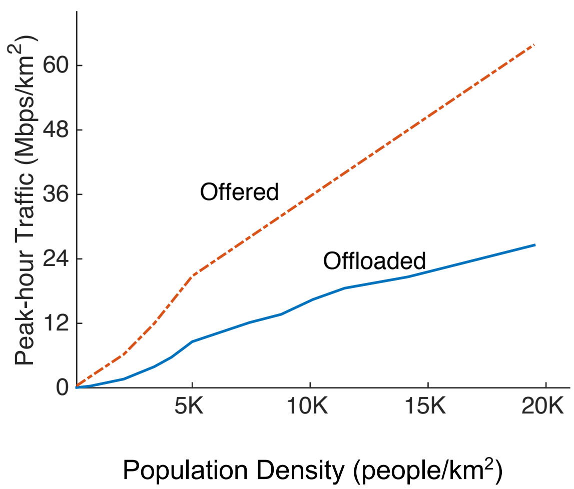

Figure 3.3. Average traffic offered and offload throughput rate at a peak hour for varying population densities and other parameters fixed at base case values, optimal RSU quantity at each point (i.e. at RSU quantity that maximizes the NPV of benefit minus the NPV of cost for each population density: 1 to 2 RSUs/km^2).

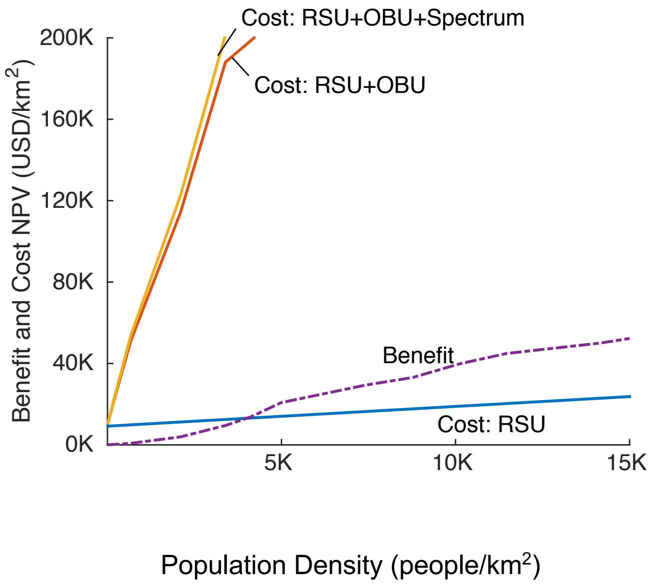

Figure 3.4. Benefit and cost for varying population densities (and other parameters at base case values), and optimal RSU quantity at each point.

Figure 3.4 shows that benefit grows faster than RSU cost for a wide range of population densities. However, the trend is different for OBU cost, which grows much faster than benefit. The reason can be seen from Figure 3.5, which shows vehicle ownership and vehicle usage as a function of population density, using base case assumptions of vehicle ownership per capita, time on the road per vehicle and peak hour ratio. Ownership refers to the total number of vehicles available per unit of area. In the event of a mandate, ownership determines how many vehicles will have OBUs installed, and the total OBU cost. On the other hand, vehicle usage is the number of vehicles on the road at peak hours. A vehicle equipped with an OBU will only have traffic offloaded carried while on the road, and only the peak-hour offload throughput is relevant for benefit. Thus, Figure 3.5 helps explain why OBU costs are significantly higher than offload benefit of Internet access, under a mandate scenario. Over locations with increasing population densities, and under uniform OBU penetration, vehicle ownership rises faster than vehicle usage, making OBU costs rise faster than benefits, at least for base case values of the other parameters. This may not be true for all assumptions. For example, if OBUs cost less, then total OBU cost would grow more slowly with respect to population density, but the OBU costs would need to be lower that baseline by one order of magnitude for OBU costs not to increase faster than benefits.

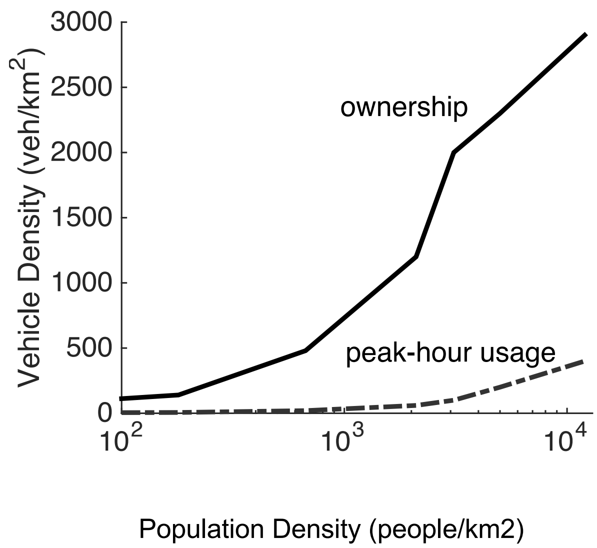

Figure 3.5. Average vehicle ownership and usage (i.e. vehicles on the road at peak hours) for population density ranges.

On the other hand, Figure 3.6 shows what happens if the ratio between the quantity of OBU-equipped vehicles in use and the total quantity of vehicles owned is different than in the base case assumption. For this graph, the population density is held in the base case value (5,000 people/km2), as well as penetration, traffic per vehicle, unit costs, spectrum characteristics and number of vehicles owned. What is varied is the number of vehicles on the road at peak hour per capita, meaning the ratio between that and the number of vehicles owned changes. The ratio value of 0.08 corresponds to the base case, which vary among cities with comparable population densities, in part due to factors like the availability of public transportation (European Commission 2012). As Figure 3.6 shows, the net benefit of deploying RSUs will be greater in a city where a larger fraction of cars is on the road in peak hours.

If vehicles are equipped voluntarily rather than because of a mandate, then Figure 3.6 is relevant for a different reason. If adoption is voluntary, owners of vehicles that are often in use are more likely to adopt, and this would also have the effect of increasing the ratio of OBU-equipped vehicles on the road at peak hour to total cars that is shown in Figure 3.6. Therefore, if many of the OBU-equipped cars are driven extensively, as is certainly the case for the OBU-equipped vehicles in Porto, then this will also increase the net benefit of deploying RSUs.

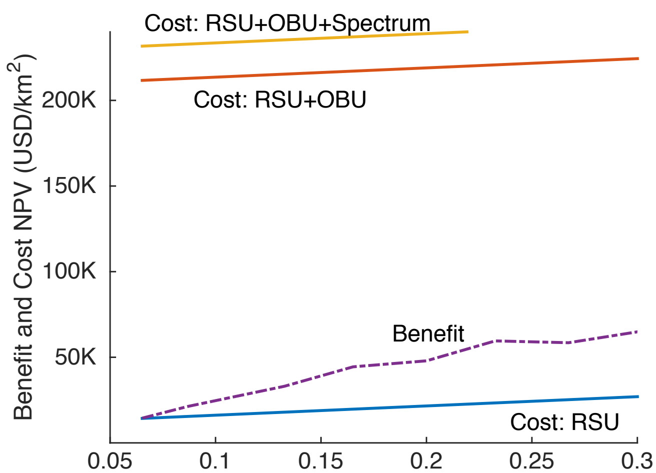

Number of Vehicles Using DSRC at Peak-hour /Number of Vehicles Owned

Figure 3.6. Benefit and cost for varying ratios between the quantity of OBU-equipped vehicles in use and the total quantity of vehicles owned (and other parameters at base case values), and optimal RSU quantity at each point (1 to 2 RSUs/km^2).

3.4.3. Impact of OBU penetration

Like population density, OBU penetration is likely to affect benefit and costs, although unlike population, penetration may increase rapidly over time. With higher penetration, it is expected that both the number of OBU-equipped vehicles and the number of OBU-equipped vehicles on the road at peak hours will increase. Therefore, it is expected that more RSUs for those vehicles to connect to will be necessary. On one hand, this makes costs of OBU and RSUs increase with penetration. On the other hand, offload throughput per unit of area from cellular to V2X, and hence benefit, are also expected to increase. This Section examines the effect of penetration on benefit and cost.

Figure 3.7 shows offload throughput as a function of OBU penetration, assuming the population density, quantity of vehicles, traffic per vehicle, unit costs and spectrum parameters are held constant at the base case values for all values of penetration considered. The graph shows the rate of traffic offered increases rapidly as a function of penetration, which is expected considering an increasing quantity of vehicles for higher penetrations. The offload throughput is also higher. If penetration increases over time as expected (especially if there is a mandate), then offload throughput will increase over time. Since benefit is defined as a function of offload throughput, RSUs are expected to be deployed only in areas where the potential rates are high enough for benefit to exceed RSU cost, as long as spectrum and OBU costs are sunk. Therefore, the growth of offload throughput over time would eventually cause the potential benefit of offloading Internet access to exceed the cost of RSUs in regions where this is not initially the case.

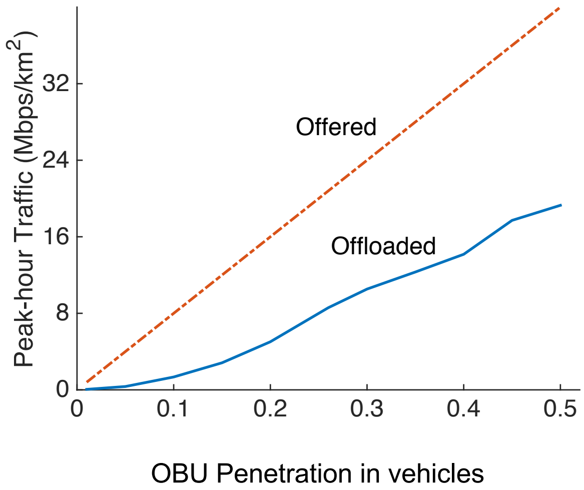

Figure 3.7. Average throughput traffic offloaded (and offered rates of incoming traffic) at a peak hour: varying OBU penetration (and other parameters at base case values), optimal RSU quantity at each point (0.8 to 1.6 RSUs/km^2).

With higher penetration, offered load per km^2, throughput, and ultimately benefit increase. Moreover, the optimal number of RSUs to carry that throughput increases with penetration as well. Figure 3.8 shows benefit and costs as a function of penetration, with all parameters, except penetration, at base case values and the RSUs densities chosen as described in subsection 3.4.2. The top horizontal axis shows penetration for a lower population density (2,000 people/km^2), while the bottom horizontal axis shows penetration for the base case population density (5,000). Figure 3.8 shows that as OBU penetration increases, benefit increases faster than RSU cost. Thus, in cities where RSU deployment does not result in benefit exceeding RSU cost within the current planning horizon, this may change after a few years as penetration increases. For the base case assumptions, the benefit of Internet access exceeds RSU costs when penetration is 0.19 or greater in a city with population density of 5,000/km^2. For a population density of 2,000 people/km^2, benefit exceeds cost only when penetration is 0.37 or greater.

However, OBU cost increases much faster than benefit, thus if penetration increases over time, the difference between OBU cost and benefit is also likely to increase. In this situation, if there were no benefits other than Internet access, then social welfare would decrease. But that could only be true if V2X communications had no safety benefits whatsoever, which is unlikely.

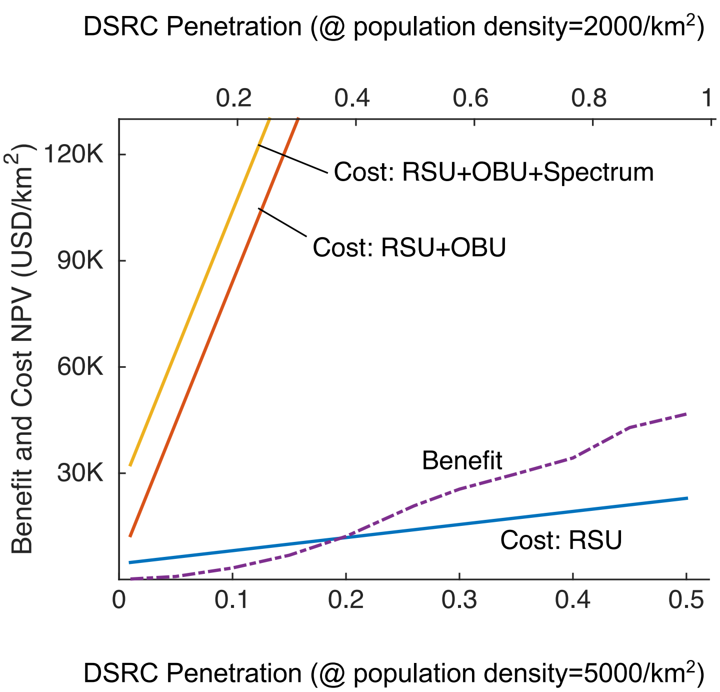

Figure 3.8. Benefit and cost for varying values of OBU penetration (and other parameters at base case values), and optimal RSU quantity at each point. Each horizontal axis refers to a different population density.

3.4.4. Impact of cost per onboard unit (OBU)

In order to investigate whether offload benefit of Internet access through V2X-based networks would ever exceed all costs, including the OBU cost that dominated in the base case, this subsection examines the effect of the OBU unit cost on total benefit and cost.

Figure 3.9 shows benefit and costs as a function of OBU unit cost, for the base case values of population density, the quantity of vehicles, penetration, traffic, RSU and macrocellular unit costs, and spectrum parameters. The quantity of RSUs is chosen to maximize the difference between benefit and RSU cost. If a mandate was to be justified by offload Internet access only, then offload benefit of Internet access alone should exceed all DSRC costs. Figure 3.9 shows that total OBU cost would exceed RSU and spectrum costs combined under these assumptions, and that the sum of RSU and OBU costs would exceed offload benefit of Internet access even if the cost per OBU falls by more than 80% from $350 to $50.

It is only possible for the cost per OBU to decrease over the range shown in Figure 3.9 if OBUs are mass-deployed at a scale comparable to Wi-Fi devices. In the physical level DSRC is specified by the IEEE 802.11p standard (IEEE 2010b), which is mostly an adaptation of the Wi-Fi 802.11a standard for the 5.9 GHz band. Wi-Fi radios with antennas currently cost no more than a few tens of dollars, and perhaps DSRC OBU costs could drop if it is mass produced. But even if this happens, Figure 3.9 shows that benefit still does not exceed total OBU cost for the base case scenario.

However, if there is a mandate in which spectrum is already allocated and OBUs are purchased, then spectrum and OBU costs are sunk. In this scenario, since offload benefit of Internet access exceeds RSU cost for base case assumptions, RSU deployment for offload Internet access does increase social welfare.

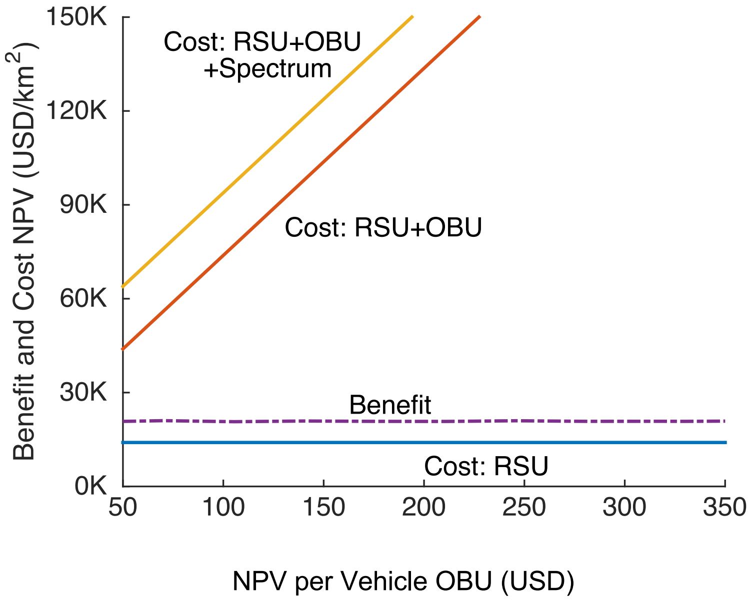

Figure 3.9. Benefit and cost for varying NPV per OBU (and other parameters at base case values), and optimal RSU quantity (1 RSU/km^2).

3.4.5. Impact of data rates of Internet traffic

It is important to consider different values for data rate per vehicle, both because there are uncertainties in any prediction of future data rate, and because data rate is generally expected to increase rapidly over time (Clarke 2014; Cisco 2015). This subsection examines the effect of traffic per vehicle on benefit and cost.

Figure 3.10 shows offload throughput as a function of traffic per vehicle, assuming the population density, quantity of vehicles, penetration, unit costs and spectrum parameters are held constant in the base-case values for all values of traffic considered. For an increase in the traffic per vehicle, offload throughput increases, though with a decreasing marginal gain. Figure

3.10 suggests that offload throughput is still growing for traffic per vehicle as high as four times the base-case value. If traffic per vehicle increases over time, then offload throughput is likely to increase over time as well even if traffic grows as much the wide range shown in Figure 3.10, under base-case values for the other assumptions.

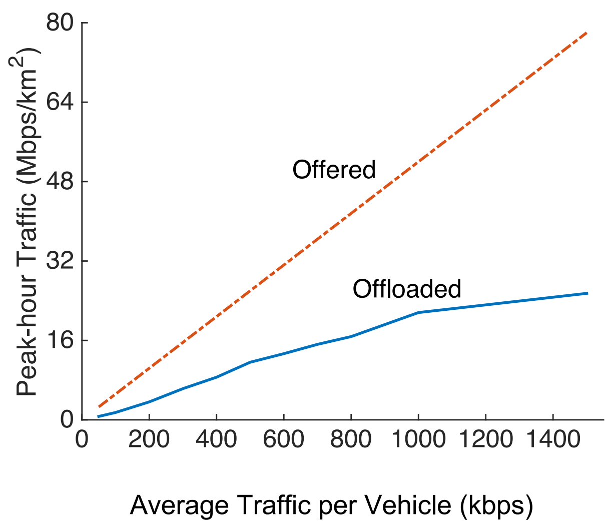

Figure 3.10. Average traffic offered and offload rate at a peak hour for varying traffic per OBU-equipped vehicle on the road, optimal RSU quantity at each point (0.9 to 2 RSUs/km^2).

Figure 3.11 shows the benefits and costs as a function of incoming traffic per vehicle, assuming the population density, quantity of vehicles, penetration, unit costs, and spectrum parameters are held constant at the base case values for all values of traffic considered, and the RSUs densities are chosen as described in subsection 3.4.2.

The difference between the benefit of Internet access and RSU cost increases with the rate of incoming traffic per vehicle, similarly as with OBU penetration. If traffic or penetration increase over time as predicted, then benefit will eventually exceed RSU cost in less populated areas where this is not the case soon after the mandate is effective. In subsection 3.4.2 it is shown that benefit exceeds RSU cost for locations with population density above 4,000 people per km^2, with the base case assumption of traffic per vehicle. Since Figure 3.11 shows that the difference between the benefit of Internet access and RSU cost increases with traffic per vehicle, and if traffic will increase over time as some predict, then benefit would exceed RSU cost in locations with population densities below 4,000 people per km^2 over time.

Figure 3.11 also shows that, under the base case scenario for the other assumptions, the benefit of Internet access exceeds RSU cost for traffic per vehicle above 250 kbps at peak hours. This corresponds to a monthly usage of 3 GB per vehicle. Thus, deploying RSUs would still result in the benefit exceeding RSU cost soon after the mandate becomes effective in the densely-populated urban area represented by our base case if data rate is about half of what some are currently predicting.

The average data rate of an OBU-equipped vehicle may also exceed the average data rate of all vehicles if vehicle owners purchase OBUs voluntarily, rather than only in response to a mandate. The owners who adopt voluntarily would be the ones who benefit the most. If owners are charged for Internet service based on usage, then more owners of vehicles with higher volumes of Internet traffic would opt in, and average data rates could be much greater than the base case. For example, a bus company offering Internet service for passengers (such as the one in Porto) might voluntarily install OBUs as soon as RSUs are operating. This is because the bus company expects a data rate per vehicle that is well above average and carrying that traffic over a cellular network would be expensive. Thus, for a given OBU penetration rate, the benefit of Internet access will exceed costs at a lower population density if there is a significant level of voluntary adoption of OBUs.

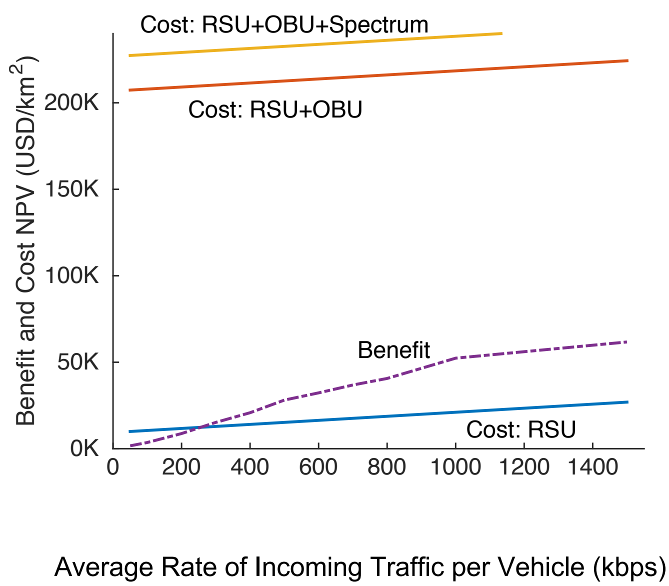

Figure 3.11. Benefit and cost for varying rates of incoming traffic per DSRC-equipped vehicle on the road (and other parameters at base case values), and optimal RSU quantity at each point.

3.4.6 Throughput under high OBU penetration or data rates

We also examined the impact on cost-effectiveness of high data rates and OBU penetration. Since benefit is proportional to throughput, we investigated whether the latter increases or collapses for high network load. As data rate of incoming traffic increases, throughput increases rapidly until it reaches a peak, and then remains within a small percentage of its peak for higher loads, regardless of RSU density. This limit at arbitrarily high load is called saturation throughput (Bianchi 2000; Chhaya and Gupta 1997), and Figure 3.12 shows the relationship between saturation throughput per km^2 and density of OBU-equipped vehicles. Data rates of incoming traffic at each path are high enough to keep the TCP transmission buffers constantly full, and curves for several RSU densities are shown. The graph shows that saturation throughput increases linearly when vehicle density is low, and then remains close to its maximum for all OBU densities above some threshold, regardless of RSU density. The fact that throughput is close to peak even for much higher loads than the base case value means that congestion and interference never cause a serious loss of throughput (and therefore benefit), probably thanks to mechanisms such as MAC-level collision avoidance and transport-layer congestion control. This is important, because the number of OBU-equipped vehicles will increase over time if the U.S. Dept. of Transportation mandates the technology for all new cars, and data rates of incoming traffic are also expected to increase sharply over time (Sandvine 2014; Cisco 2015). As a result, cities with vehicular networks need not fear that benefit will decline as load goes up every year.

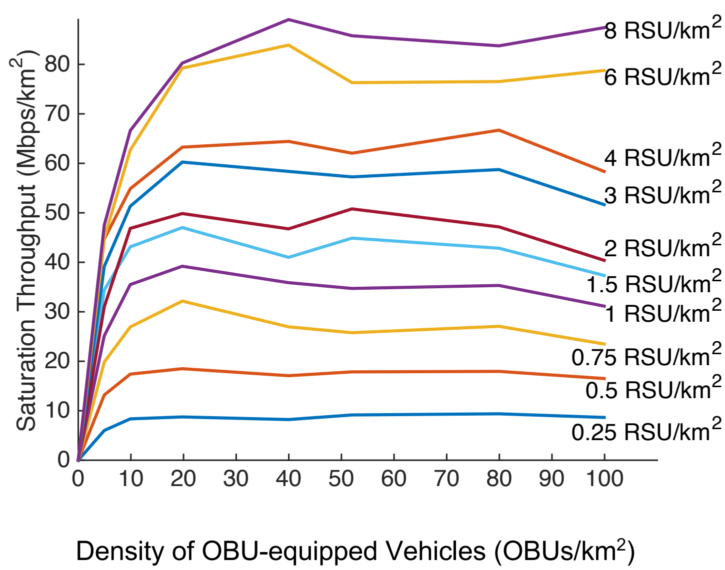

Figure 3.12. Saturation throughput of the vehicular network for varying density of OBU-equipped vehicles and data rates of 100 Mbps/OBU. Each line refers to a fixed RSU density.

3.4.7. Impact of cost per roadside unit (RSU)

If the cost per RSU is lower than in the base case, then it may be worthwhile to deploy more RSUs to increase total throughput. On the other hand, if RSUs are significantly more expensive than in the base case, then that may prevent deployment and result in no benefit at all. This Section examines the effect of RSU unit costs on total benefit and cost.

Figure 3.13 and Figure 3.14 show throughput, benefit and costs as a function of RSU unit cost. The base case values of population density, the quantity of vehicles, penetration, traffic, OBU and macrocellular unit costs, and spectrum parameters are assumed. The quantity of infrastructure for each value of RSU unit cost is chosen to maximize the difference between benefit and RSU cost, as explained further in 3.4.2. The cost per RSU affects that optimal quantity of RSUs, which influences V2X-based throughput. This is shown in Figure 3.13: if the cost per RSU is lower than the base case value ($14,000), then offload throughput is higher and vice versa. However, even with that variation in throughput, Figure 3.14 shows that the total benefit and cost results are robust to a wide variation of costs per RSU. Even if this cost is 30% higher (or lower) than the base case, benefit of Internet access will still exceed total RSU cost.

However, that result might change if the cost per RSU is radically different than the base case. For example, if RSUs are deployed by businesses in places that require expensive poles or lack of access to commercial power or communications, then the cost per RSU might be much higher, and Figure 3.14 shows that benefit of Internet access is lower when total RSU cost if its unit cost is higher than $20,000 per RSU and other assumptions are at base case. On the other hand, if the decision to deploy RSUs are made by a municipality that already has pole, energy and backhaul infrastructure available, cost per RSU may be low, and RSU deployment might be beneficial even for less densely populated cities than the “threshold” density shown in section 3.4.2 for base case assumptions, as long as spectrum and OBU costs are sunk under a mandate.

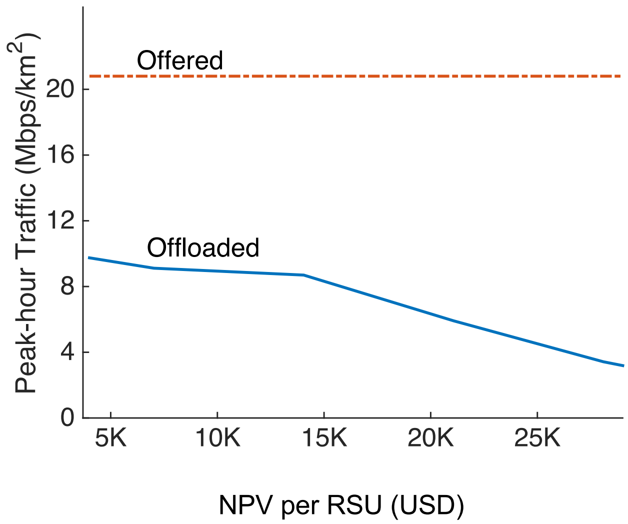

Figure 3.13. Average offload rate at a peak hour for varying PV per RSU, and optimal RSU quantity at each point (1.3 to 0.8 RSU/km^2).

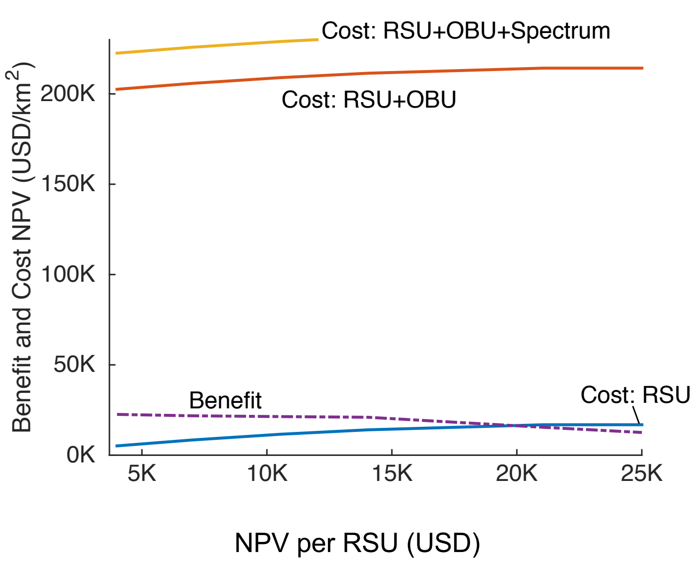

Figure 3.14. Benefit and cost for varying PV per RSU (and other parameters at base case values), and optimal RSU quantity at each point.

3.4.8. Impact of macrocellular factors

Macrocellular costs are expected to influence benefit and costs in the opposite way as the V2X-related costs analyzed in the previous subsections. Figure 3.15 shows benefit and costs as a function of the unit cost per macrocellular tower, for which is assumed the base case values of population density, the quantity of vehicles, penetration, traffic, V2X costs and spectrum parameters. If the NPV of the cost per macrocellular tower is higher than the base case assumption, then benefit of Internet access exceeds RSU cost in less populated areas than in the base case scenario. On the other hand, if macrocellular cost is lower than in the base case, than the benefit might be lower than in the base case. However, Figure 3.15 shows that the findings in previous subsections do not change substantially if the cost per macrocellular tower changes over a range of 20% below or above the base case value.

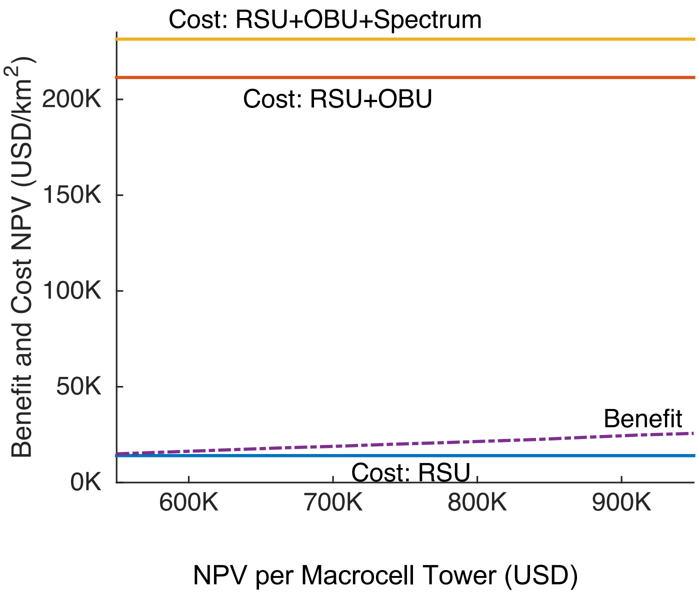

Figure 3.15. Benefit and cost for varying NPV per macrocellular tower, and optimal RSU quantity (1 RSU/km^2).

Benefits and costs may also be influenced if the amount of spectrum a carrier has available varies from the base value. Figure 3.16 shows benefit and costs as a function of the bandwidth available per carrier, and base case values of population density, quantity of vehicles, penetration, traffic, unit costs and ITS spectrum, and indicates that benefit of Internet access exceeds RSU cost if as much as 20% more bandwidth per carrier is in use. Spectrum holdings for cellular service may increase over time as long as the growing demand for mobile Internet trigger decisions to reallocate spectrum from other uses to cellular, or new technologies enable use of bands that were previously unusable for cellular. However, spectrum reallocations are not frequent and take years to become effective – 65 MHz were auctioned in 2015, being the first significant addition to mobile spectrum since 2008 in the U.S. (Clarke 2014). Therefore, over a given period the amount of spectrum may increase less than the rapid growth expected for traffic per vehicle (which increases benefit, as shown in subsection 3.4.5), which suggests that the growth in cellular spectrum may still result in Internet access over V2X-based networks being cost-effective.

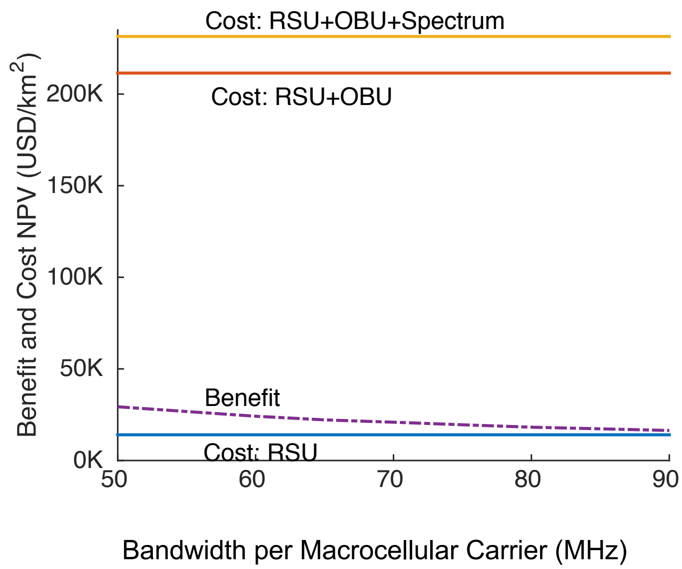

Figure 3.16. Benefit and cost for varying bandwidth available for macrocells, and optimal RSU quantity at each point (1 RSU/km^2).

3.5. Conclusions

In this chapter, we analyze the cost-effectiveness of Internet access through V2X-based networks, when compared to mobile Internet traffic being carried instead over microcellular networks. We examine throughput, benefits and costs with a given amount of ITS spectrum being used to carry Internet traffic, and with an ISP deploying RSUs exclusively for Internet access. We find that if there has already been a mandate to deploy OBUs in new vehicles, then the deployment of RSUs for Internet access increases social welfare. This is true for dense urban areas, when OBU penetration is representative of a few years after a mandate becomes effective, peak-hour Internet traffic per vehicle is compatible with forecasts for the next years, and even when those RSUs are not shared with safety or other applications. Moreover, RSU deployment is likely to become welfare enhancing in the future for many less-populated areas as well, as long as penetration or Internet traffic increases over time.

Under a mandate to deploy OBUs, our results show that the OBU cost alone exceeds benefit. However, it has been estimated that an OBU mandate will accrue significant road safety benefits (U.S. Department of Transportation 2016; Harding et al. 2014), which has motivated the allocation of ITS spectrum and the possibility of a mandate to deploy OBUs in all new vehicles in the U.S. If this mandate occurs, then the decision of whether to use V2X-based networks for Internet access becomes a decision about whether to deploy roadside infrastructure that can serve as a gateway to the Internet. For this decision, both OBU and spectrum costs would be sunk, and if the benefit of Internet access exceeds RSU cost, then a decision to deploy RSU infrastructure would increase social welfare. Our results show that benefit does exceed RSU cost under base case assumptions, which correspond to dense urban areas.

Benefits and costs are both affected by population density. If all else is equal, the benefit of Internet access through V2X-based networks minus the cost of RSUs is greater when population density is greater. With base case assumptions, benefit exceeds RSU cost in locations with population density above 4,000 people per km^2, i.e. only in fairly densely populated urban areas. However, this should change over time. Under an OBU mandate, the volume of traffic per vehicle and OBU penetration are both likely to rise rapidly beyond our baseline assumptions in the coming years. With this growth, our results show that the benefit of Internet access minus RSU costs also increases. Thus, if all assumptions are close to base case values except OBU penetration and traffic per vehicle, then the benefit will exceed the cost of RSUs in regions with lower and lower population densities over time. Therefore, the deployment of RSUs will become social-welfare-enhancing over more of the country. However, there will remain areas where deployment of RSUs does not enhance social welfare, including those rural areas where population density is so low that cellular networks are not capacity-limited, i.e., they have excess capacity and don't need offload.

Since benefit is proportional to throughput, we also examined how it scales for high levels of load in the vehicular network, which is likely to happen in the future. For example, the base-case data rates assumed for vehicles may be a small fraction of total Internet traffic today. However, both that fraction and the overall traffic may increase over time when more of certain types of vehicles are connected, such as buses and self-driving cars where passengers access the Internet rather than drive. We find that even for arbitrarily high loads, throughput per unit of area (and thus benefit) approaches a saturation level that remains close to the maximum achievable throughput, meaning that the cost-effectiveness of vehicular networks will not decline even as Internet traffic and the penetration of OBUs in vehicles grow sharply as predicted.

RSU cost also affects whether deployment of RSUs would increase social welfare, and RSU cost varies from community to community. For example, all else being equal, benefits of Internet access through DSRC minus RSU costs will be lower where the provider has to acquire infrastructure (poles, backhaul, etc.), than where RSUs are deployed by a municipality that already has infrastructure available, or where part of the RSU cost is incurred for another purpose, e.g. a given RSU is shared for safety and Internet traffic.

Like any model of a complex system, our analysis is based on a number of simplifying assumptions, some of which may be explored further in future research, such as the variability in traffic per vehicle and among vehicle types, and the dynamics of traffic, penetration, and costs over time. However, the conclusion that the benefit exceeds RSU cost in urban areas but is lower than the sum of RSU, spectrum, and OBU costs is sufficiently robust, such that a small change of around 20% in any of these assumptions would not change it. If reality differs from the base case even more than this, this is most likely either because of our assumption about a mandate or our assumption about mobile traffic levels. For example, if data rates are substantially higher (or lower) than our baselines estimate of 5 GB per month per vehicle, then the population density required for the benefit of Internet access through V2X to exceed the cost of RSUs may be less (or more) than our estimated 4,000 people per km^2, respectively.

Chapter 4. Sharing Roadside Infrastructure for the Internet of Vehicles

4.1. Introduction

In the previous chapter, we have shown that deploying RSU infrastructure is cost-effective to offload mobile Internet traffic from cellular onto V2X-based mesh networks. This applies for an entity such as a municipality or an Internet Service Provider (ISP) deploying RSUs exclusively for Internet access. However, governments will likely deploy RSUs for other purposes as well. In particular, V2X communications have the potential to improve road safety, and V2X-enabled OBUs may be mandated for all new cars (U.S. Department of Transportation 2016). However, substantial investment in RSUs is required, which was estimated to be in the order of billions of dollars nationwide (Wright et al. 2014). RSUs for safety probably will not be deployed until state and local governments choose to pay, which may slow adoption of V2X applications.

This chapter is about cost savings from infrastructure sharing, when it is deployed by government agencies and shared with private parties. If RSU cost can be reduced, V2X-related safety benefits may be experienced sooner by more people. We show that governments might save by sharing safety RSUs with ISPs for a fee.

Moreover, governments may widely deploy other types of infrastructure that could be shared. As illustrated in Figure 4.1, one example is the deployment of “smart” streetlights with communications capability, to aid services such as surveillance, air quality monitoring, etc. Those streetlights may be opportunities for ISPs of cheap access to power, poles and backhaul, and possibly available in more locations than safety RSUs. In this chapter, we also consider sharing of streetlights.

By sharing safety RSUs or streetlights, governments might charge prices to maximize either government savings or social welfare. The contributions of this chapter are to determine the prices the government would charge an ISP to achieve either goal. We consider the scenario where the amount of ITS spectrum is that allocated in 1999 (75 MHz), and the U.S. Department of Transportation (DOT) mandates vehicles to be equipped with DSRC-capable OBUs as proposed in 2016. In-vehicle Internet is increasing sharply, and ISPs must decide whether to expand cellular capacity or to deploy RSUs to offload part of the demand. These RSUs can either be deployed for Internet only by the ISP or shared. In this scenario, the ISP pays to share government infrastructure. However, the results are also applicable to some other sharing arrangements, such as a public-private joint deployment.

We analyze government infrastructure expenses, ISP infrastructure expenses, and government revenues from ISPs. We estimate these without sharing, and with sharing as a function of the price government charges to share an RSU. We assume that ISPs design their systems to carry a given volume of traffic, and ISPs minimize cost by choosing any combination of deploying their own RSUs that serve as Internet gateways, sharing safety RSUs or smart streetlights with government for a fee, and deploying traditional macrocells.

We estimate that government could save about one-fifth of the nationwide cost of safety RSUs in the U.S. if they are shared with ISPs. We also estimate an increase in social welfare from sharing safety RSUs. In the case of sharing smart streetlights, we find that nationwide benefits could be up to one-third higher than with sharing of safety RSUs. The prices that maximize government savings and social welfare may differ. However, we find that maximizing government savings results in near-optimal social welfare. The benefits of sharing would increase significantly if the Internet traffic or OBU penetration grow over time.

One aspect of our method is an engineering-economic model to estimate RSU costs, government revenues from ISPs, and the resulting government savings and increased social welfare from sharing. Some of these costs depend on how much traffic can be offloaded from macrocells to a vehicular network as a function of RSU quantity. Offload capacity is estimated with the simulation model combined with data from Portugal, as described in Chapter 2. We have extended the engineering-economic approach to address government savings and social welfare when safety RSUs or smart streetlights are shared with ISPs for a fee.

Part of the findings discussed in this Chapter was presented at the IEEE Vehicular Technology Conference (Ligo, Peha, and Barros 2017), and is published in the IEEE Transactions on Intelligent Transportation Systems journal (Ligo and Peha 2018a).

This chapter is organized as follows. Section 4.2 discusses related work, while Section 4.3 describes the data used and how our engineering-economic approach has been extended to address the issues presented in this chapter. Results are discussed, and sensitivity analysis is performed in Section 4.4. Section 4.5 concludes the chapter.

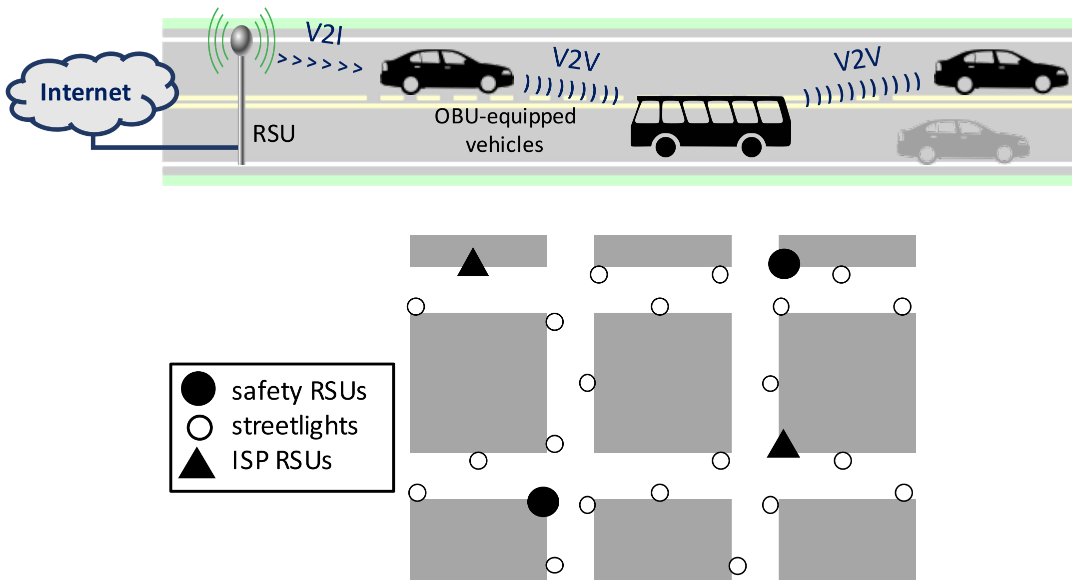

Figure 4.1. Representation of a vehicular connection to an Internet-connected RSU. An ISP may deploy its own RSUs, and it may use safety RSUs or smart streetlights shared by the government.

4.2. Related work

Although government agencies often deploy infrastructure only for their own use (Daher and Vinel 2013), previous work has shown other instances where government can save by sharing infrastructure with commercial companies. For example, as shown in (Hallahan and Peha 2011a, 2010; Peha 2013a; Hallahan and Peha 2011b), a highly cost-effective way to provide communications for emergency responders involves sharing infrastructure between government and commercial cellular providers. This approach was adopted in FirstNet, a nationwide network for emergency responders which Congress funded in 2012 with $7 billion (Peha 2013a).

Similarly, governments might share RSU infrastructure with ISPs. Some claim that demand for mobile Internet will grow sharply (Sandvine 2016). That includes in-vehicle Internet access, which is currently served mainly by macrocells that would continuously need expansion where networks are capacity-limited. Although that extra capacity is costly, previous work has shown that vehicular networks could provide Internet access at a lower cost than cellular networks. For example, it has been shown (Lu et al. 2013) that roadside microcells provide Internet access at a lower cost than cellular networks, assuming greenfield deployment of either infrastructure. In Chapter 3, we have shown that ISPs can provide Internet access at lower cost using V2X-based networks than through expanding cellular infrastructure in some regions, if ISPs deploy RSUs that function as Internet gateways. If ISPs could use government RSUs for less than the cost of their own RSUs, then ISPs might offer V2X-based Internet in more locations. Thus, there is benefit in sharing dual-use RSUs for both safety and Internet access. To the best of our knowledge, this work is the first that quantifies that benefit.

4.3. Method and system model

We have used the engineering-economic approach presented in Chapter 2 to model throughput and costs of the V2X-based network, and the cost savings from reducing the expansion of macrocellular networks when part of the growth in mobile Internet traffic is offloaded from macrocells onto V2X-based networks. We have extended that engineering-economic approach to assess government savings and social welfare when RSU infrastructure is shared, and to inform what prices should be charged for sharing. The data used and the modeling of costs, ISP strategy, government savings and social welfare from sharing are described below.

4.3.1. *Dataset*

To determine how sharing affects government savings and social welfare, for the analysis in this chapter we use the data from the real DSRC network operating in Porto, Portugal, as described in Chapter 2. More specifically, we use the GPS positions from the dataset to determine the positions of the vehicles in the simulation of throughput of the vehicular network.

In addition, we use coordinates of urban road intersections. The municipality of Porto made available the latitude and longitude coordinates from all of the city's 4,900+ road intersections. We use these coordinates of intersections to determine the locations of safety RSUs as described below.

4.3.2. *Costs of V2X and cellular infrastructure*

As in Chapter 3, we consider the amount of ITS spectrum allocated in 1999, and there is a mandate to equip cars with OBUs for safety, as may occur in the U.S. (U.S. Department of Transportation 2016). In this scenario, spectrum and OBU costs are incurred for safety and only RSU costs matter for non-safety purposes.

In our model, when Internet traffic is carried over the V2X-based network at peak hours, fewer macrocellular towers are needed than in a scenario without V2X. The cost of those fewer towers is defined in Equation 2.9 (Chapter 2). From a social welfare perspective, if that avoided cost exceeds the cost of the V2X network, it does not matter who benefits from the avoided cost, such as the ISP or its users in the form of reduced retail prices of Internet service.

However, in this chapter we are interested both in social welfare and in government savings. Because of the latter, we assess the distribution of cost savings between ISPs and governments.

We assume in this chapter that if the avoided cost of macrocells exceeds the cost of the V2X network for the ISP, then this difference is a profit for the ISP. Otherwise, the ISP is better off by not deploying V2X infrastructure for Internet access. On the other hand, deployment costs for the ISP are affected by whether RSUs are shared by the government, and at what price. Therefore, if RSU sharing reduces infrastructure cost for the ISP, then its profit is higher than in the absence of sharing. We assume the ISP will adopt the RSU deployment strategy that maximizes profit. Also, the amount of Internet traffic does not depend on whether it is carried over macrocells or RSUs (shared or not). Thus, ISP revenue does not depend on strategy, so the ISP strategy that maximizes profit also minimizes cost. If this strategy includes shared RSUs, then government savings and increased social welfare are possible.

All costs are defined as the sum of upfront and ongoing costs over the base case time horizon, which are discounted to present values using the base case discount rate presented in Chapter 2.

4.3.3. Locations of safety RSUs and smart streetlights

The ISP strategy in deploying RSUs, government savings and social welfare from sharing depends on the quantity and locations of safety RSUs or streetlights that can be shared. For the former we assume the density of safety RSUs that can be shared is

$$0.2\,P\,/\,1000 \tag{4.1}$$

which is based on (Wright et al. 2014), (FHWA 2016), where P is the population density. Safety RSUs are placed at the intersections with the highest average quantity of vehicles at peak hours. This assumption is consistent with (Wright et al. 2014), which found that a significant number of crashes are intersection-related and high-volume intersections are likely to have the highest number of crashes. We also assume that placement and quantity of safety RSUs do not depend on whether they are shared. The locations of safety RSUs in the simulation model are based on the real intersection coordinates from Porto.

We also examine sharing of smart streetlights, which we assume can be upgraded to provide V2X-based Internet access and are ubiquitous. Therefore, they are available at the locations that would be chosen by an ISP deploying its own RSUs (intersections or not). We name the density of locations that can be shared (either safety RSUs or streetlights) as N_{sa}.

4.3.4. ISP strategy for using shared and Internet-only RSUs

Cost for the ISP per unit of area (km^2) is

$$C_{ISP} = p\, N_{sh} + c_{io}\, N_{io} \tag{4.2}$$

where p is the price per shared RSU, c_{io} is the cost the ISP bears to deploy an Internet-only RSU by its own, and N_{sh} and N_{io} are the densities of shared RSUs and Internet-only RSUs per km^2 that maximize ISP profit $(AC_{tower} - C_{ISP})$. Note that N_{sh} and N_{io} affect not only C_{ISP} but also the avoided cost of macrocells AC_{tower}, because N_{sh} and N_{io} affect throughput.

For sharing of safety RSUs in a given scenario of population density, OBU penetration and other assumptions, we find the N_{sh} and N_{io} that maximize $(AC_{tower} - C_{ISP})$ according to the following procedure. We run the simulation with the density of RSUs $N_{sh} + N_{io}$ ranging from 0 to

10 RSUs/km2. and with the density of shared RSUs ranging from 0 to *min*{density of safety RSUs, $N_{sh} + N_{io}$}. For each density, we calculate throughput and costs, and thereby determine the optimal N_{sh} and N_{io}.

For each RSU density, RSUs are initially placed where they are likely to result in the most throughput. Thus, RSUs should be set in places with a large number of vehicle positions at peak hours. More specifically, RSUs are placed using the *k-means* clustering heuristic as described in Chapter 2. If all RSUs are Internet-only, then the RSUs remain at these locations. For cases where some RSUs are shared, RSUs are moved to be collocated with safety RSUs until the desired density of shared RSUs is reached. If *j* RSUs are to be moved, then we move the RSUs that are closest to an unshared safety RSU. For the case of sharing of smart streetlights, the locations of shared RSUs are the same as the Internet-only RSUs, because streetlights are assumed to be ubiquitous.

4.3.5. Social welfare and government savings from sharing

As outlined in Section 4.1, governments might choose prices for sharing RSUs that maximize either social welfare or government savings. We assume that both total Internet traffic carried and the availability of safety-enhancing applications do not depend on the number of RSUs deployed or shared, so consumer benefit is not affected by RSU strategy. As a result, social welfare is maximized by carrying that traffic and supporting those safety applications with the combination of RSUs and macrocells that result in the lowest overall cost. In contrast, if governments choose to maximize government savings, they would seek to collect as much as possible from ISPs, without considering how RSU strategy might benefit Internet users and providers. As a result, governments that maximize government savings may deploy a different number of shared RSUs and share RSUs at different prices from those governments that maximize social welfare.

In our model, social welfare is increased when V2X-based Internet access is provided at a lower cost than using macrocells for vehicular users. The increase in social welfare when there is no sharing SW_n (NPV per km^2) is given by

$$SW_n = AC_{tower,n} - C_n \tag{4.3}$$

where $AC_{tower,n}$ is the avoided cost of macrocells (NPV per km^2) under the ISP strategy that maximizes $(AC_{tower,n} - C_{ISP})$, calculated with Equation 2.9, in the absence of sharing. C_n is the cost (NPV per km^2) of Internet-only RSUs that would be deployed in the absence of sharing, and given by

$$C_n = c_{io}\, N_{io}^{nosharing} \tag{4.4}$$

where $N_{io}^{nosharing}$ is the density of Internet-only RSUs deployed when there is no sharing.

The increase in welfare under sharing is

$$SW_{sh} = AC_{tower,sh} - C_u - C_{io} \tag{4.5}$$

where $AC_{tower,sh}$is the avoided cost of macrocells (NPV per km2) calculated with (2) when RSUs can be shared, and $C_{io} = c_{io}\, N_{io}^{sharing}$ is the cost to deploy $N_{io}^{sharing}$ Internet-only RSUs in the sharing case. C_u is the cost to upgrade safety RSUs or streetlights for sharing, per km^2. C_u is defined as

$$C_u = c_u\, N_{sh} \tag{4.6}$$

where c_u is the cost to share a safety RSU or streetlight.

Sharing results in a net increase in social welfare if and only if the increase under sharing SW_{sh} exceeds the increase when there is no sharing SW_n. The net increase (NPV per km^2) is

$$SW = SW_{sh} - SW_n \tag{4.7}$$

The price p affects the density of RSUs N_{sh}, which affects social welfare. The lower p, the greater is N_{sh}. However, if p is lower than the cost to share c_u, then the ISP will deploy RSUs which marginal $AC_{tower,sh}$ is lower than their marginal cost, and this decreases social welfare. To find the pricing strategy that maximizes social welfare, we differentiate Equation 4.5 with respect to N_{sh}:

$$\frac{\partial SW_{sh}}{\partial N_{sh}} = \frac{\partial AC_{tower,sh}}{\partial N_{sh}} - \frac{\partial C_u}{\partial N_{sh}} - \frac{\partial C_{io}}{\partial N_{sh}} = \frac{\partial AC_{tower,sh}}{\partial N_{sh}} - c_u \tag{4.8}$$

as long as the variation of C_{io} with respect to N_{sh} is negligible. From the above, SW_{sh} is maximized when $\frac{\partial AC_{tower,sh}}{\partial N_{sh}} = c_u$. Since the ISP will deploy shared RSUs as long as $\frac{\partial AC_{tower,sh}}{\partial N_{sh}} \geq p$ (i.e. the macrocell cost avoided by an additional RSU exceeds the price to the ISP), then SW_{sh} (and SW) is maximized when $p = c_u$.

Government savings from sharing is

$$GS = (p - c_u)\, N_{sh} \tag{4.9}$$

The price p that maximizes GS is not obvious, because Equation 4.9 depends on N_{sh}, which is also affected by p.

Besides, a positive GS results in a secondary effect. Each dollar of GS means that a dollar less is required from public funds (raised from taxes) to finance safety RSUs or streetlights. Taxation causes a social burden known as the excess burden of taxation, which has been estimated to be between 0.3 and 0.5 of public funds raised (Triest 1990). If government savings means less taxes, then the excess burden is also reduced. We call this reduction an "avoided" excess burden, or AEB. We assume a positive GS causes an AEB of

$$AEB = 0.4\, GS \tag{4.10}$$

4.3.6. Base-case scenario

The base case numerical values for the assumptions described in Chapter 2 are also used for the analysis in this chapter. Those include the cost of each RSU deployed by the ISP.

For the unit cost to share a safety RSU or smart streetlight for Internet access (c_u), we assume that the average NPV is $1,400. This cost is assumed as the incremental cost of backhaul on safety RSUs is streetlights. In (Clark, Lehr, and Bauer 2011) the backhaul cost is estimated as about $1 per Mbps per month. The NPV results from incurring costs for 16 Mbps of capacity. (The throughput per RSU is below 16 Mbps in more than 95% of the simulations.)

Densities of safety RSUs or smart streetlights are derived as described in subsection 4.3.3. In the base scenario, we consider RSU sharing with ISPs. However, the method applies to any provider of IP-based traffic that would typically be carried over macrocells, such as mobility and environmental applications (Wright et al. 2014).

4.4. Results and discussion

In this section we show the RSU deployment strategy that maximizes ISP profit, the pricing strategies of a government that seeks to maximize either social welfare SW or savings GS when charging a profit-maximizing ISP for shared RSUs, and the national implications of those government strategies. Moreover, we perform sensitivity analysis to show the impact of the most important assumptions on nationwide results.

4.4.1. ISP strategy for using shared and Internet-only RSUs

In this subsection we discuss the ISP strategy, i.e. the densities of shared RSUs N_{sh} and Internet-only RSUs N_{io} that maximize the ISP profit from RSU deployment $(AC_{tower} - C_{ISP})$.

First, we found that throughput of a shared safety RSU is less than 5% different from the throughput of an Internet-only RSU for 95% of them. This is shown in Figure 4.2. Thus, if an Internet-only RSU is cost-effective in a location, and there is a safety RSU or streetlight available for sharing nearby, then the ISP will use the shared RSU as long as $p < c_{io}$ (i.e. the price of sharing is lower than the cost of an Internet-only RSU).

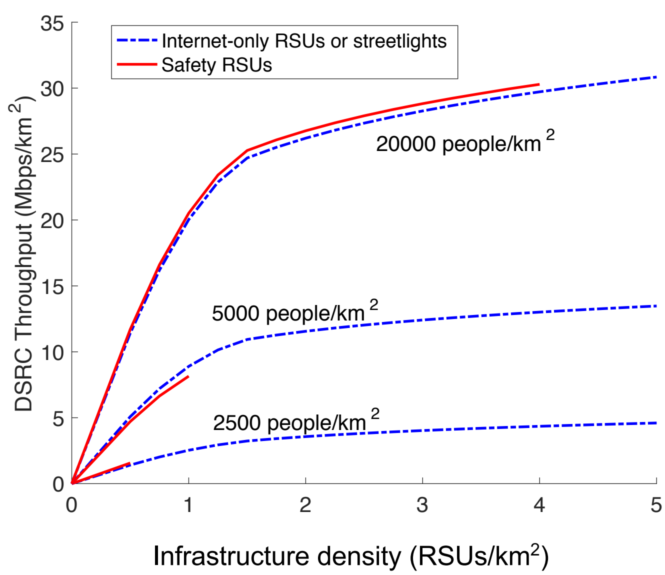

Figure 4.2. Throughput as a function of RSU density, for different population densities. The dashed lines show throughput from Internet-only RSUs, which is the same as the throughput of RSUs located at smart streetlights, while the solid lines show throughput of

Internet data through sharing of safety RSUs. There are less safety RSUs than Internet-only RSUs because it is assumed that there are 0.2 safety RSUs per 1,000 people.

We also found that the ISP strategy is affected by conditions that vary with population density. That is, there is a different strategy under each of three mutually-exclusive conditions, defined by the RSU densities N_{sa} and $N_{io}^{nosharing}$ (see subsection 4.3.5). We label those conditions I, II and III, as shown in Figure 4.3.

Condition I is $N_{io}^{nosharing} = 0$, i.e. in the absence of sharing the ISP strategy is to not deploy Internet-only RSUs. However, if the price of shared RSUs is lower than the avoided cost of macrocells, then the ISP deploys a density of shared RSUs N_{sh}.

Condition II is $N_{sa} > N_{io}^{nosharing} > 0$. For a price lower than the avoided cost of macrocells, the ISP strategy is to use more RSUs than it would deploy without sharing ($N_{io}^{nosharing}$).

Condition III is $N_{io}^{nosharing} \geq N_{sa} > 0$, i.e. the density of Internet-only RSUs $N_{io}^{nosharing}$ that maximizes ISP profit under no sharing is higher than the density of shareable locations. In that case, an ISP would profit from deploying $N_{io}^{nosharing}$, but there are not as many shareable locations as the ISP would deploy. Thus, the ISP strategy is to use all shared RSUs as long as $p < c_{io}$. Also, the ISP may deploy Internet-only RSUs in locations not served by safety RSUs or smart streetlights.

Figure 4.3 (a) shows N_{sa} for safety RSUs and $N_{io}^{nosharing}$, both as a function of population density. The graph shows that $N_{sa} > N_{io}^{nosharing}$ (i.e. condition I or II) for most population densities. Condition I applies for population densities below 4,000 people/km^2, while condition II applies for most populations above that density. However, there is a narrow range of population densities around 5,000 people/km^2 where condition III holds. On the other hand, Figure 4.3 (b) shows that the density of smart streetlights will always exceed $N_{io}^{nosharing}$, thus there is no population density where condition III applies.

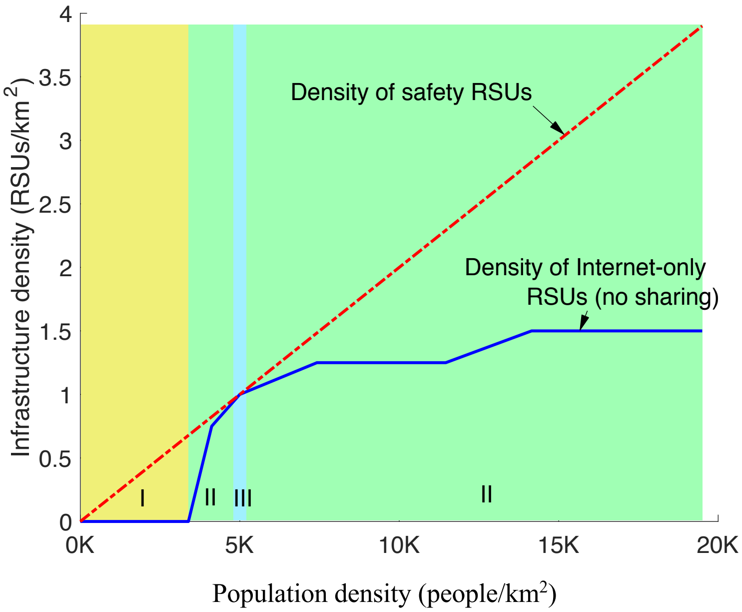

(a) Safety RSUs

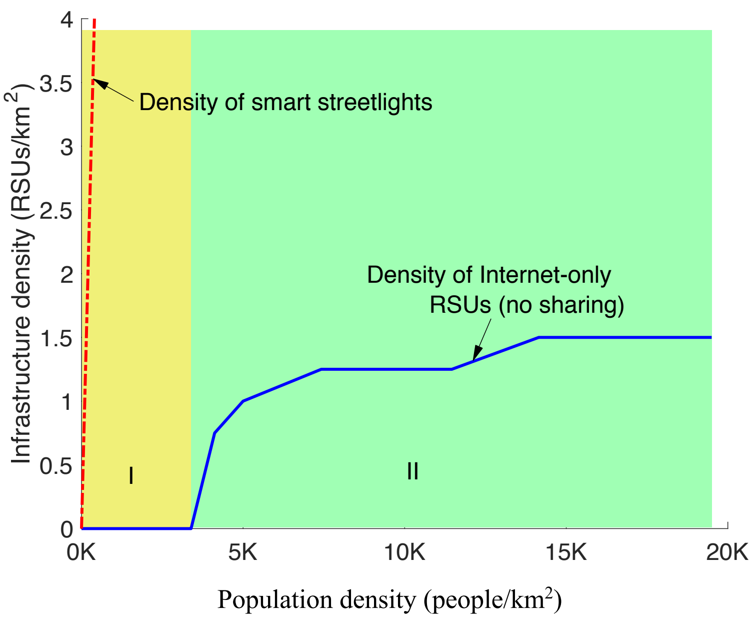

(b) Smart Streetlights

Figure 4.3. RSU density as a function of population density. The solid line is the density $N_{io}^{nosharing}$. The dashed line is the density N_{sa}. The background colors represent which condition (I, II or III) applies for each population density.

4.4.2. Government strategy to maximize social welfare SW

This Section discusses the pricing strategy that maximizes social welfare from sharing. In Subsection 4.3.5 we show that SW is maximized by setting price $p = c_u$. (Since c_{io} = \$14,000 and c_u = \$1,400, the optimal p/c_{io} is 0.1.) Figure 4.4 (a) shows that for sharing of safety RSUs, SW is maximized for $p = c_u$, but remains at its maximum for other prices as well. This is because there is a range of prices where all safety RSUs are shared.

For population densities where condition I holds, there is a limit for the price p above which SW is zero. This is because no RSUs are deployed at p near c_{io}, since the avoided cost of macrocells is below RSU cost. The curve for 2,500 people/km^2 illustrates one population density under condition I. For condition II SW is maximum for $p = c_u$, but then falls with p. This is shown for 20,000 people/km^2. For safety RSUs (Figure 4.4 a) SW is maximum for $p/c_{io} = 0.1$ (i.e. $p = c_u$) and for higher prices sharing and SW decrease. For condition III, if $p < c_{io}$, all safety RSUs are shared and SW is maximum. This is illustrated in Figure 4.4 (a) for 5,000 people/km^2.

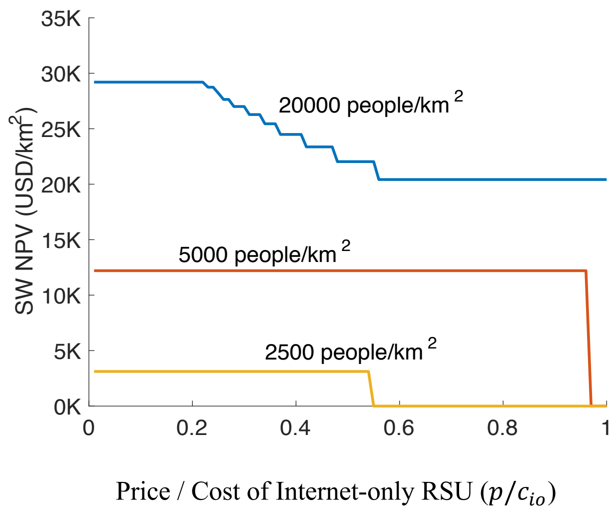

(a) Sharing of safety RSUs

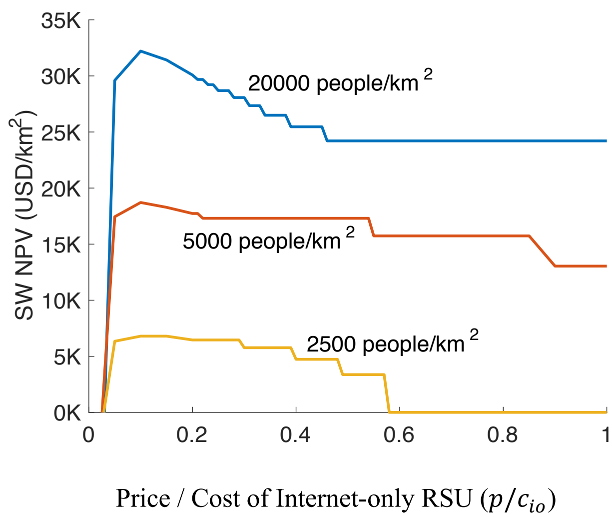

(b) Sharing of smart Streetlights

Figure 4.4. 10-year NPV per km^2 of social welfare from sharing SW as a function of price for sharing, for different population densities.

For streetlights, Figure 4.4 (b) shows that SW is maximized for $p = c_u$ and decreases for other prices, although there is still a range where SW is close to maximum. Moreover, the maximum SW from sharing streetlights, as shown in Figure 4.4 (b), is higher than the maximum

SW from sharing safety RSUs in Figure 4.4 (a). This is because there are less safety RSUs than the quantity the ISP would use at the optimal price. That relative gain in SW from sharing streetlights instead of safety RSUs is larger for lower population densities than for higher population densities, because of the diminishing incremental benefit per additional RSU. For example, the maximum SW from sharing streetlights is twice that from sharing safety RSUs at 2,500 people/km^2, while that gain is only 10% higher at 20,000 people/km^2.

In summary, a government seeking to maximize SW can set $p = c_u$ under all conditions. However, the magnitude of SW shown in Figure 4.4 may differ for assumptions other than those considered. For example, as discussed in Chapter 2 we believe that OBU cost does not change much when more radios are used for Internet access, when compared to less radios. If this cost difference is otherwise high, then OBUs will likely be deployed with fewer radios, resulting in a somewhat lower SW.

4.4.3. Pricing strategy to maximize government savings GS

The sharing price p determines how much of the cost saving from sharing RSUs increases either ISP profit or GS.

In areas where condition I holds, there is a price limit above which $GS = 0$. Figure 4.5 illustrates that for 2,500 people/km^2. The government would charge p/c_{io} of about 0.5 for maximum savings. For condition II, a large quantity of shared RSUs are deployed at a low price, but fewer shared RSUs are used as they become more expensive for the ISP. For 20,000 people/km^2 GS is maximized by setting p/c_{io} close to 1 in the case of sharing safety RSUs (Figure 4.5 a). This is also true for streetlights (Figure 4.5 b, see 5,000 and 20,000 people/ km^2). For condition III (5,000 people/km^2 in Figure 4.5 a), all safety RSUs are shared as long as $p <$

c_{io}. In this case, a government would again charge p close to c_{io}. In any case (I, II or III), adopting a price strategy of charging the maximum price the ISP can bear is optimal.

The GS resulting from charging the maximum price the ISP can bear is similar between sharing of safety RSUs and sharing of smart streetlights. This is because at the maximum price the ISP can bear, the ISP is going to deploy the same quantity of shared RSUs regardless of type (safety RSUs or streetlights).

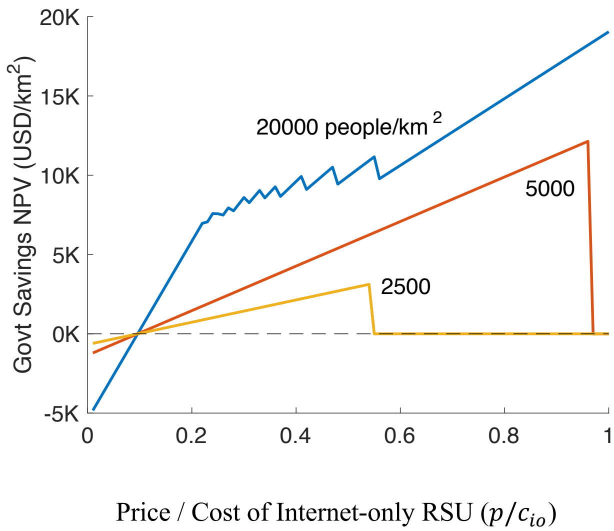

(a) Sharing of safety RSUs

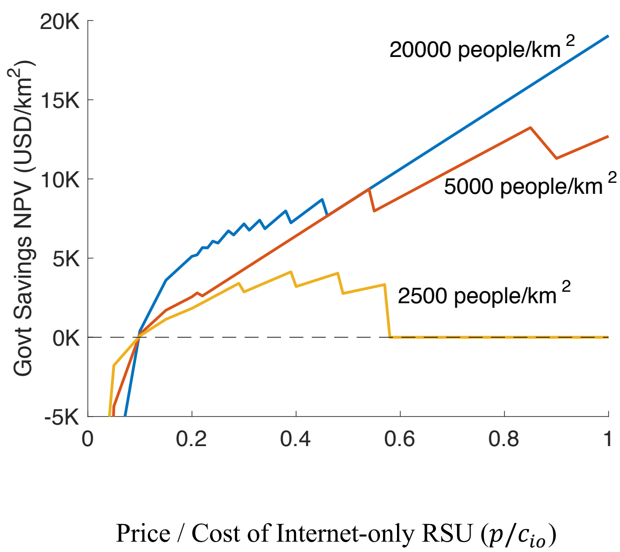

(b) Sharing of smart Streetlights

Figure 4.5. 10-year NPV per km^2 of government savings from sharing GS as a function of price, for different population densities.

Also, in Figure 4 6 we show that for sharing of safety RSUs at locations with densities around 5,000 people/km^2 (condition III), the ratio between government savings and the total cost of safety RSUs can be over 80%, because most or all RSUs can be shared at a high price. However, for other population densities, the ratio is lower because of the price limits discussed above. For higher population densities such that the quantity of safety RSUs is higher than the optimal number of Internet-only RSUs, the safety locations with less Internet benefit are not used.

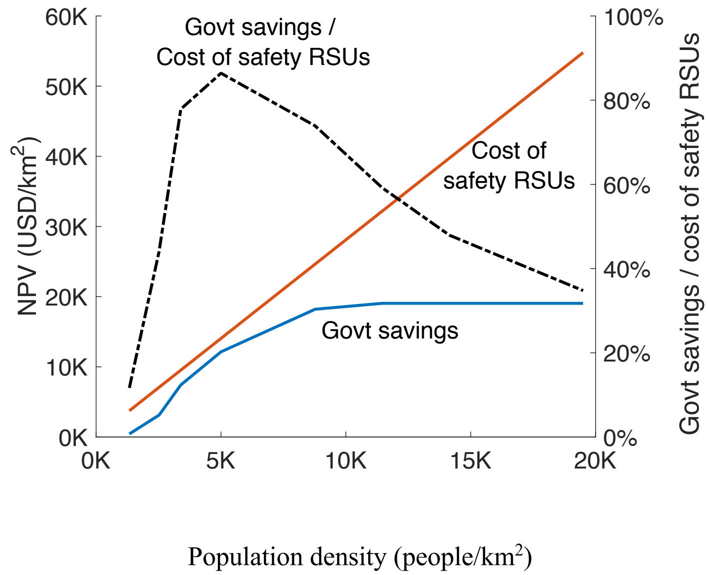

Figure 4.6. 10-year NPV per km^2 (left axis) of government savings from sharing, and the cost of safety RSUs. The right axis refers the ratio between savings and cost of safety, shown in the dashed line.

4.4.4. Government trade-offs and avoided excess burden AEB

In many regions, government savings GS and social welfare SW cannot be maximized at the same price. While $p = c_u$ is optimal for SW, the p that maximizes GS varies with population density. Therefore, there is a trade-off between maximizing SW and maximizing GS for some population densities.

One way to reconcile the two objectives is to consider avoided excess burden (AEB), as defined in subsection 4.3.5. Thus, aside from the objectives of maximizing GS or SW, a third possible objective for the government might be to maximize $SW + AEB$, a hybrid objective that depends on both GS and SW.

Figure 4.7 shows $SW + AEB$ as a function of price, for different population densities. The graph shows that $SW + AEB$ does not always increase monotonically with price p. The pricing strategy that maximizes $SW + AEB$ depends on population density. However, Figure 4.7 suggests that charging the maximum price the ISP can bear is near optimal, i.e. the $SW + AEB$ obtained with such a strategy is not more than 10 or 20% lower than the maximum $SW + AEB$. Thus, a strategy of maximizing GS is similar to maximizing $SW + AEB$. Moreover, $SW + AEB$ from sharing of streetlights (Figure 4.7 b) is higher than $SW + AEB$ from sharing of safety RSUs (Figure 4.7 a).

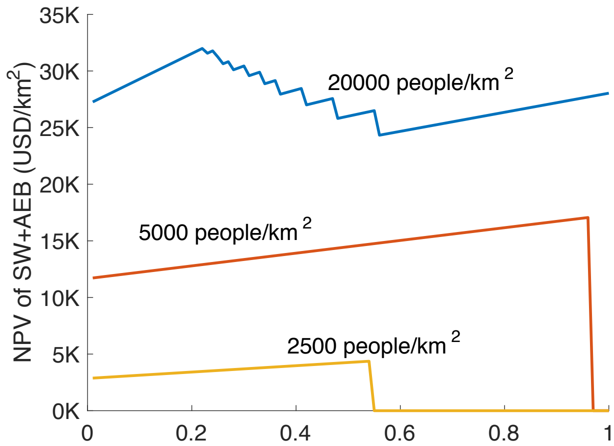

Price / Cost of Internet-only RSU (p/c_{io})

(a) Sharing of safety RSUs

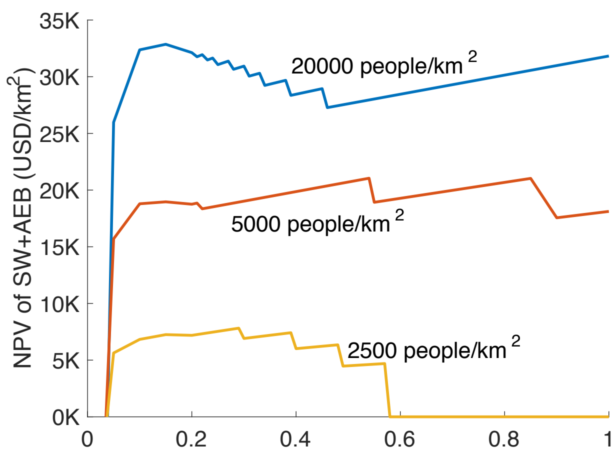

Price / Cost of Internet-only RSU (p/c_{io})

(b) Sharing of smart Streetlights

Figure 4.7. 10-year NPV per km^2 of social welfare plus the avoided excess burden (SW+AEB) as a function of price, for different population densities.

4.4.5. *Nationwide government savings and social welfare*

In this subsection, we quantify the nationwide effects of RSU sharing. We assume the population density variation of the U.S., and that all census tracts determine their pricing strategies to either maximize social welfare SW, maximize government savings GS, or maximize SW plus avoided excess burden AEB.

GS, SW and AEB were calculated for each U.S. census tract, using 2010 data from (United States Census Bureau 2015), then summed nationwide. Penetration, data rates and other assumptions are fixed in the base values. The results are shown in Figure 4.8. For sharing of safety RSUs, Figure 4.8 (a) shows that the 10-year NPV of nationwide GS is close to $200 million when the pricing strategy is to maximize GS. Assuming (i) there are about 310 thousand signalized intersections in the U.S. and safety RSUs would be deployed in about 20% of those intersections ((Wright et al. 2014), Table 7) in the period of analysis, and (ii) a safety RSU has the same cost c_{io} of an Internet-only RSU, then the cost of nationwide deployment of safety RSUs would be about $310000 \times 0.2 \times \$14000 = \$850$ million. Thus, Internet access could save about $200 million / $850 million = 23% of the investments in safety RSUs by local governments.

On the other hand, Figure 4.8 (a) shows that nationwide $SW + AEB$ for sharing of safety RSUs is just 2% lower when maximizing GS is the objective, compared to $SW + AEB$ when the objective is to maximize $SW + AEB$. Thus, if state/local governments lean to the objective of maximizing GS, the nationwide impact in $SW + AEB$ seems to be small.

Figure 4.8 (b) shows nationwide results for smart streetlights. The graph shows that the maximum NPV of nationwide SW and $SW + AEB$ are higher than the nationwide results with sharing of safety RSUs, which indicates the advantage of having more locations that can be shared in the streetlight case. For example, nationwide $SW + AEB$ with the price strategy to maximize GS is \$270 million from sharing of safety RSUs and \$360 million from sharing of streetlights, or 33% higher than the former. This is because the density of shared RSUs N_{sh} is higher for streetlights than for safety RSUs, especially when price is low such as in locations under condition I.

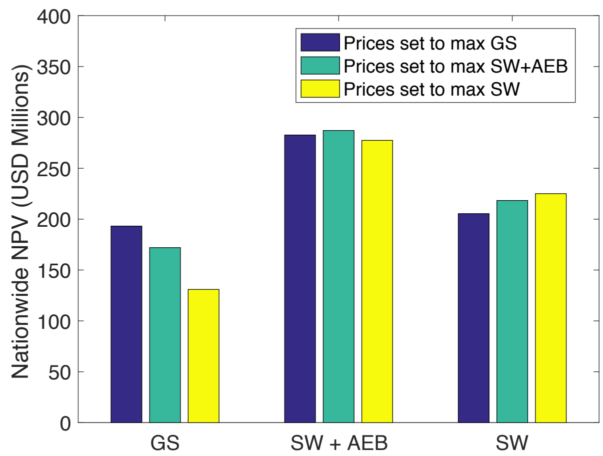

(a) Sharing of safety RSUs

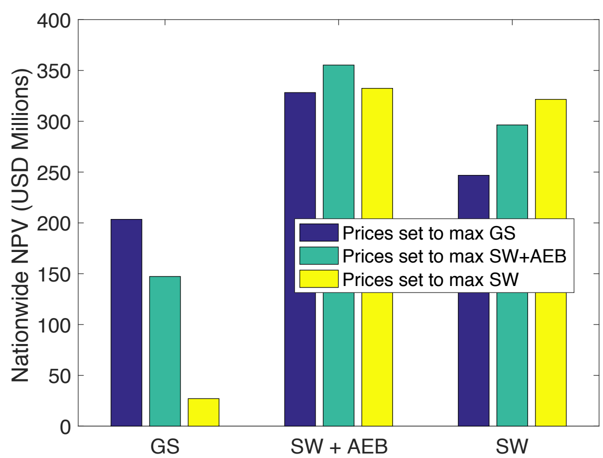

(b) Sharing of smart Streetlights

Figure 4.8. 10-year NPV, summed over U.S. census tracts, of GS, SW, and SW+AEB. Prices are chosen at each census tract to maximize GS (blue bars), SW+AEB (green), or SW only (yellow).

4.4.6. Sensitivity analysis

The results presented above depend on the numerical assumptions presented in Chapter 2 and in Section 4.3. Some of those assumptions are expected to increase over time, such as Internet data rates and OBU penetration in vehicles. Other assumptions are uncertain, such as the costs c_{io}, c_u, and of macrocells. This subsection investigates the robustness of the results with respect to the assumptions that are most likely to vary, are most uncertain or have the most impact.

Figure 4.9 shows the effects of variations (one assumption at a time) on the nationwide social welfare plus the avoided excess burden $SW + AEB$. The variations are shown for safety RSUs in Figure 4.9 (a) and for smart streetlights in Figure 4.9 (b). The graphs show that data rate per OBU has the highest effect on nationwide $SW + AEB$ from sharing of either safety RSUs or streetlights. The reasons are twofold. First, we considered a variation for data rates that is higher than the variation for the other assumptions. This is because it has been reported that the volume of mobile Internet traffic has grown 70% per year (Cisco 2015), and thus estimates of data rates over multiple years are uncertain. On the other hand, it is also uncertain whether the current growth in mobile Internet will hold in the future for vehicular users. Hence, we consider variations of up to twice and down to half the base data rate in Figure 4.9. The second reason for the high impact of data rates on results is that higher rates both raise savings GS and welfare from sharing SW in a location. Moreover, data rate determines the number of RSUs to deploy (shared and not shared). There are locations where V2X-based Internet is not cost-effective at the base data rate, but eventually become cost-effective as data rates increase. A consequence is that the variation in nationwide $SW + AEB$ is more than proportional to the variation in data rate per OBU-equipped vehicle. For example, Figure 4.9 shows that if data rates are twice the base rate, nationwide $SW + AEB$ is 7 times the base value for sharing of safety RSUs and 18 times for streetlights.

That also explains why varying the penetration of OBUs in vehicles has a significant impact. For sharing of safety RSUs an increase of 25% in penetration results in an increase of 20% in nationwide $SW + AEB$. However, we considered a variation in penetration much smaller than the variation in data rates because the growth in the former is expected to be relatively low, even in the case of a mandate. For example, the US DOT estimated that penetration would reach 50% in no less than 10 years (U.S. Department of Transportation 2016).

Uncertainty may also have a major impact. Regarding the cost of macrocells, the more expensive is the cost of a tower, the higher is the benefit of Internet over shared RSUs. For example, land and legal costs can be major components, which vary by location. Hence, we consider a variation of plus or minus 50%. If a macrocell costs on average half of the base assumption, then Figure 4.9 shows a high reduction in $SW + AEB$, although $SW + AEB$ is still greater than zero. That would mean less savings and a smaller increase in social welfare than predicted with base assumptions, and V2X-based Internet might be cost-effective in fewer locations than predicted in the base case scenario.

The uncertainty in the other factors seems to have limited effect on nationwide results. Regarding c_u, even if it is 50% higher than the base value, the variation in nationwide $SW + AEB$ is less than 20% either for sharing of safety RSUs or streetlights. This is partly because we believe c_u is relatively small compared to c_{io}, and hence the nationwide results should be robust to the uncertainty in c_u. The uncertainty on the cost of an Internet-only RSU c_{io} should be high, because deployment can be cheap in locations with mounting structure, energy and backhaul available, while c_{io} can be much higher than the base value in locations with no such infrastructure. Figure 4.9 (a) shows that 25% cheaper Internet-only RSUs cause a roughly proportional decrease in $SW + AEB$, because the optimal price to share is near c_{io} for a wide range of population densities. However, Figure 4.9 (b) shows that variations of 25% in c_{io} have negligible effect on nationwide results for streetlights.

For data rates or OBU penetration higher than the base values, and at low sharing prices, one may conclude (wrongly) that benefit exceeds the cost for the ISP and trigger deployment of shared RSUs even for population densities close to zero. Actually, cellular networks in sparsely populated areas are likely to be coverage-limited instead of capacity-limited, implying no benefit of offload. For this chapter, we assumed that benefit is zero for population densities below 10 people/km^2. This is reasonable because for a random sample of U.S. counties, those with population densities below 10 people/km^2 have shown average cell radius of tens of km, while most counties with more than 10 people/km^2 have lower and decreasing cell range as population density increases (which is an indication that those cells are capacity-limited).

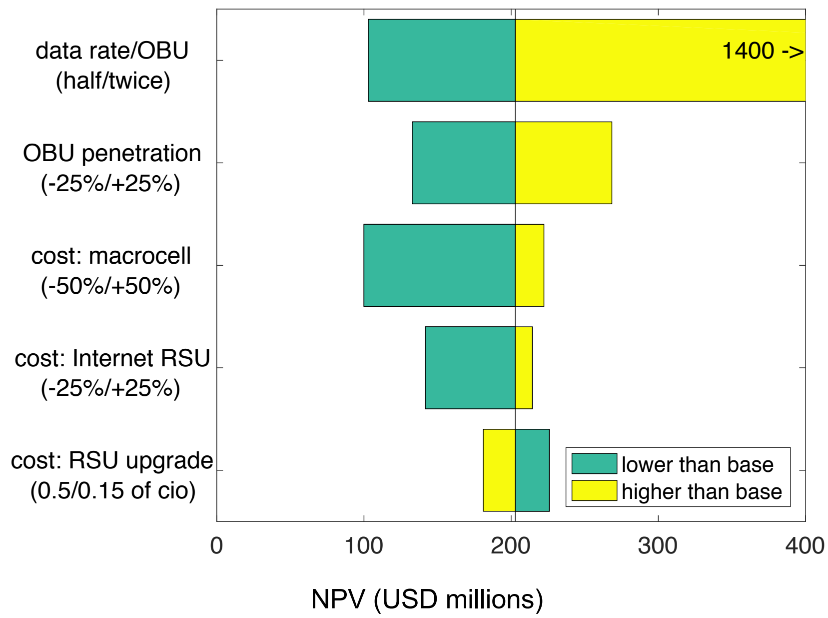

(a) Sharing of safety RSUs

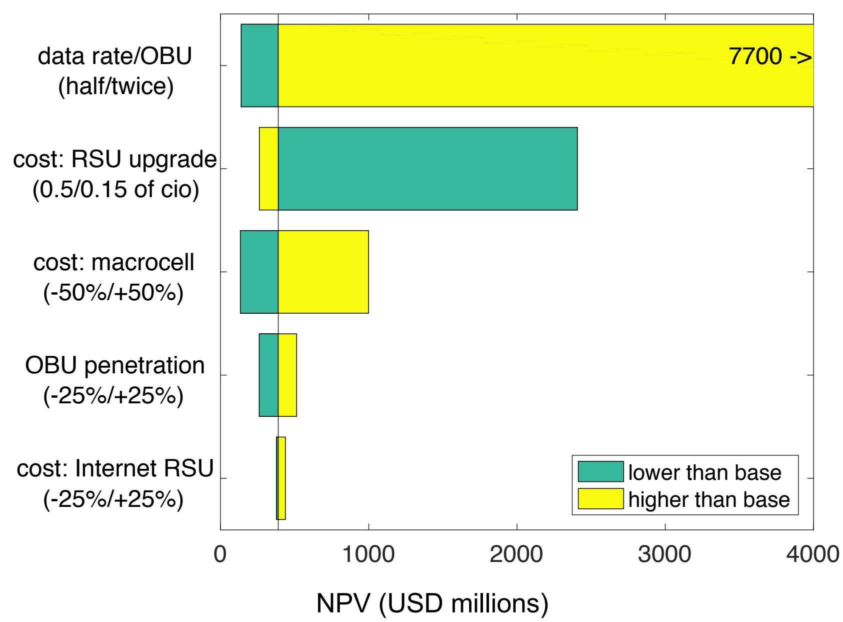

(b) Sharing of smart streetlights

Figure 4.9. 10-year NPV, summed over U.S. census tracts, of social welfare from sharing SW plus the avoided excess burden of taxation AEB. Prices are chosen at each census tract to maximize SW+AEB. The vertical line in each graph is the nationwide result with the assumptions in base values. Each horizontal column refers to a variation in one of the numerical

assumptions (data rate per OBU, OBU penetration, c_{io} or c_u), and the values in parentheses indicate the range of variation in the assumption.

4.5. Conclusion

In this chapter we assess cost savings from sharing of infrastructure for vehicular communications, when it is deployed by government agencies and shared with private parties. We show that sharing RSUs deployed for safety or smart streetlights with ISPs would result in savings for the government who owns them, and these savings could be used to offset investment. Sharing would also enhance social welfare, when compared to RSUs being deployed independently by ISPs for Internet access only.

Moreover, we show that the pricing strategy a government should adopt to charge an ISP for sharing depends on location, with respect to population density. If price is lower than the cost of Internet-only RSUs, then an ISP is likely to deploy more RSUs with sharing than without it. In particular, shared RSUs are deployed in locations where Internet-only RSUs are not cost-effective. Thus, sharing allows V2X-based Internet over more areas of the country than it would be the case without sharing.

Government savings from sharing safety RSUs or smart streetlights are maximized when the price to share is close to the cost of Internet-only RSUs, for locations where Internet over V2X-based networks is cost-effective even without sharing. However, for places with lower population densities, there is a price above which ISPs do not deploy RSUs, so there is no revenue for the government. For a nationwide deployment, we estimate the savings as 23% of the total investment in safety RSUs. In addition, we found that maximum government savings are similar between safety RSUs and streetlights. The reason is that at the prices that maximize savings, ISPs will be indifferent between their own RSUs or shared ones, regardless the latter are safety RSUs or streetlights.

If a government chooses to maximize social welfare, the optimal price equals the cost to share RSUs. At this price, social welfare from sharing is different between sharing of safety RSUs and smart streetlights. Welfare is maximized at prices where the ISP will deploy many more shared RSUs than the ISP would deploy on its own. Because there are more streetlights than safety RSUs, SW is higher for streetlights than safety.

The pricing strategy that maximizes government savings often differs from the strategy that maximizes social welfare. However, the effect of such a trade-off in nationwide social welfare plus the avoided excess burden of taxation $SW + AEB$ is limited. If state and local governments choose to maximize savings, the resulting $SW + AEB$ is close to maximum.

Moreover, we found that nationwide $SW + AEB$ is one third higher for sharing of smart streetlights than for sharing of safety RSUs, when the price strategy is to maximize savings.

If a government chooses to maximize savings, it probably has inaccurate information about the maximum price the ISP can bear. For each location, there is a price limit above which the ISP will not deploy any shared RSU, and this limit depends on the population density of the location and costs experienced by the ISP. These costs are unknown to governments. If more than the maximum price is charged, then the ISP will choose not to share. That is why governments may choose to maximize savings and charge less than the maximum price the ISP can bear. If that happens, governments would still experience $SW + AEB$ within 20% of its maximum.

Some of the numerical assumptions adopted in this work are likely to increase over time, while others are uncertain. A sensitivity analysis revealed that cheaper macrocells may result in lower nationwide $SW + AEB$. On the other hand, if data rates or OBU penetration grow over time as expected, nationwide $SW + AEB$ increase more than proportionally to that growth. Moreover, we found that uncertainty in factors such as the cost of an Internet-only RSU and the cost to upgrade safety RSUs or streetlights have limited effect on nationwide results.

Chapter 5. Spectrum Allocation and Sharing

5.1. Introduction

The U.S Federal Communications Commission (FCC) has allocated 75 MHz of spectrum in the 5.9 GHz band for Intelligent Transportation Systems (ITS) since 1999 (Lansford, Kenney, and Ecclesine 2013). This so-called “ITS band” has been regulated for use by V2X devices operating in accordance with DSRC standards since 2004 (U.S. Federal Communications Commission 2004). In previous chapters, we found that it is cost-effective to deploy V2I infrastructure for Internet access under certain assumption. Those assumptions include the use of the ITS band as having fixed size and being used exclusively by V2X devices, as currently defined in the aforementioned FCC regulations. In this chapter we relax those conditions, since the question of whether ITS should have an exclusive allocation of 75 MHz is currently debated. For that matter we address three interrelated issues. One is how much spectrum should be made available for ITS, rather than for other purposes. Allocating spectrum for ITS has an opportunity cost, which is the foregone benefit of not using that spectrum for other purposes. If the benefit of allocating more spectrum for Internet access than the amount currently available for ITS exceeds the opportunity cost, then that extra allocation enhances social welfare. Otherwise, social welfare would be enhanced if that spectrum is made available for other uses. A related issue that we address is whether the ITS band should be shared with non-vehicular devices. In 2013, the FCC issued a Notice of Proposed Rulemaking (NPRM) to permit unlicensed devices in that band (U.S. Federal Communications Commission 2013). With such sharing the utilization of the ITS band is expected to be improved, by allowing devices mainly with Wi-Fi technology to use that spectrum. As a result, the capacity of Wi-Fi access would be expanded. However, to date there has been no consensus on the rules to be adopted for such sharing (Lansford et al. 2015), because sharing the ITS band may cause harmful interference to DSRC communications. The third issue is if the ITS band is to be shared, then what form of sharing should be implemented. In particular, we examine whether unlicensed devices should active cooperate to deliver V2X traffic in order to gain access to the ITS band.

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 examine the economic benefit of V2X-based Internet access. More specifically, we assume that a certain amount of spectrum is sufficient to serve road safety applications, and then explore whether adding spectrum would result in other benefits, such as offloading traffic from cellular onto V2X networks. In Chapter 2, we showed that deploying V2X infrastructure for offload is cost-effective in urban areas. This will likely be relevant for the foreseeable future. Although macrocellular capacity continues to increase as carriers expand infrastructure and regulators allocate more spectrum for cellular communications, mobile Internet traffic has grown 18-fold in the past 5 years (CIsco 2017), justifying alternative approaches such as data offload from cellular to other networks, which we evaluated in Chapter 2. Moreover, in Chapter 3 we have also shown that it is even more cost-effective if infrastructure is shared between Internet access and safety applications. However, those works considered the bandwidth allocated for ITS as fixed and not shared. In contrast, this chapter focuses on spectrum management; we examine the economic benefit of using spectrum to offload Internet traffic in excess of what is needed for safety communications. 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 issues of whether and how ITS spectrum should be shared with unlicensed devices, we address two other research questions. One is what the difference among throughputs to vehicles and unlicensed devices in spectrum exclusively allocated, and throughputs in shared spectrum with devices following different sharing schemes. These schemes differ on whether V2X and unlicensed devices merely coexist in shared bands by avoiding mutual interference whenever possible, or unlicensed devices rather cooperate with V2X devices by carrying V2X traffic in addition to their own. Another research question is how much spectrum is needed to carry a given amount of data from V2X 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, with the several sharing schemes considered in this chapter.

The debate over the FCC NPRM is primarily about which spectrum-sharing scheme causes less interference to safety-related communications. In contrast, the contribution of our work describe in this chapter is that 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. (This is consistent with one of the proposals to the NPRM.) In that scenario, Internet Service Providers (ISPs) deploy V2X-based roadside units (RSUs) connected to the Internet to offload part of the growing volume of vehicular traffic. In addition, ITS spectrum allocated for non-safety ITS traffic can be shared with unlicensed devices such as Wi-Fi hotspots, according to one within several possible sharing schemes.

The analysis performed in this chapter is based on the economic-engineering approach described in Chapter 2, which is extended to address the research questions presented above.

One extension is the simulation of Wi-Fi hotspots representing unlicensed devices, which coexist and/or cooperate with V2X devices on shared channels. Besides the several inputs of the simulation that are varied, in this chapter we also vary the amount of spectrum used by V2X and by unlicensed devices. The densities, locations, and data rates of both vehicles and hotspots are also varied.

Parts of the findings discussed in this chapter were presented at TPRC – Research Conference on Communications, Information and Internet Policy (Ligo and Peha 2017), and will be presented at the IEEE DySPAN – International Symposium on Dynamic Spectrum Access Networks (Ligo and Peha 2018b).

This chapter is organized as follows. Section 5.2 discusses related work, while Section 5.3 describes in detail how the engineering-economic method, data and assumptions are extended to address spectrum-related issues. Results are presented in Section 5.4, and our conclusions are in Section 5.5.

5.2. Related work

To the best of our knowledge, this is the first work that addresses all the research questions presented above, and in particular, this is the first work that examines whether it is cost-effective to allocate more or less ITS spectrum than what is currently allocated, based on economic marginal benefit and marginal cost of that spectrum for offloading Internet traffic.

The development of V2X and the allocation of the ITS band were motivated primarily by road safety applications. Hence, DSRC standards require that safety-related communications have priority over IP traffic (Kenney 2011). This requirement was taken into account in the two leading schemes proposed for the FCC NPRM on spectrum sharing (U.S. Federal Communications Commission 2013). One defining feature of a scheme is whether to allow primary-second sharing or sharing among equals. Another feature of the sharing arrangement is whether devices coexist, or rather actively cooperate to avoid mutual interference (Peha 2009). One proposal (Lansford et al. 2015) is based on primary-secondary sharing without cooperation (Peha 2009) from the primary devices, which means that licensed devices are given priority to use spectrum. Unlicensed devices are allowed to use the same spectrum when and only when their transmissions would not cause harmful interference to DSRC transmissions. In the proposal, unlicensed devices are allowed to use *all channels* of the ITS spectrum, but they must stop transmitting in a channel when any DSRC transmission is detected. Supporters of the proposal argue that it doesn't require any change in channel assignments on the ITS band nor in V2X devices, which have already been extensively tested.

The second proposal (Qualcomm 2013) 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. 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." Supporters of the coexistent sharing-among-equals proposal argue that it more effectively protects the reliability of safety-related DSRC messages, which would still have a portion of exclusive spectrum allocated. (Pang, Padden, and Alderfer 2018) found that unlicensed devices operating on bands adjacent to those used by DSRC devices do case significant interference to DSRC transmissions.

Previous research addressed several issues on the coexistence between DSRC and Wi-Fi devices in the ITS band, including (Naik, Liu, and Park 2018, 2017; J. Wang et al. 2016; Cheng et al. 2017; Liu, Naik, and Park 2017; Park and Kim 2014; Lansford, Kenney, and Ecclesine

2013; K.-H. Chang 2015). (Lansford, Kenney, and Ecclesine 2013) and (K.-H. Chang 2015) noted that today's Wi-Fi devices cannot decode DSRC preambles (since the latter operate in 10 MHz channels), and the power spectral density is somewhat less restricted in Wi-Fi standards than in DSRC. This may result in significant interference to DSRC devices, but remedies such as requiring that Wi-Fi devices increase the interval between transmitting packets above a certain threshold can mitigate such interference (Lansford, Kenney, and Ecclesine 2013). (Park and Kim 2014) used a pair of nodes to emulate 100 vehicles coexisting with two Wi-Fi interferers. Similarly as in the previous works cited, Park and Kim found that improving Wi-Fi receiver sensitivity and increasing inter-frame spacing mitigates interference to DSRC devices. (Liu, Naik, and Park 2017) developed an analytical model of a single road with simulation, which accounts more accurately for the hidden node problem, with a Wi-Fi device always with packets to send. They found that Wi-Fi devices can interfere significantly with DSRC devices, which is consistent with (Park and Kim 2014). Moreover, under certain conditions the results from Liu et al. show that Wi-Fi devices can have negligible impact on DSRC performance. As with other works, this one concludes that improving receiver sensitivity inter-frame spacing in Wi-Fi devices can nearly eliminate interference to DSRC devices. (Cheng et al. 2017) compared different "detect and vacate" proposals, which are similar to the one proposed to the FCC by Cisco. The authors noted that coexistence between DSRC and Wi-Fi devices may be subject to the "delayed detection" problem, i.e. Wi-Fi devices take longer to detect DSRC transmissions if the Wi-Fi devices are transmitting themselves. However, the authors found that this problem is more significant under high Wi-Fi load, which might not necessarily be the case for all unlicensed devices in range of interference to vehicles.

While most of the papers above focused on the impact of Wi-Fi on DSRC safety messages, some exceptions (Naik, Liu, and Park 2017; J. Wang et al. 2016) also examined how the coexistence of DSRC devices affect the performance of Wi-Fi devices. (Naik, Liu, and Park

2017) performed experiments with one pair of DSRC devices and one pair of Wi-Fi devices, and found that "certain channelization options, particularly the high-bandwidth ones, cannot be used by Wi-Fi devices without causing interference to the DSRC nodes." However, similar to other works, this one found that improving receiver sensitivity and increasing inter-frame spacing in Wi-Fi transmissions can nearly eliminate interference to DSRC devices, without compromising Wi-Fi performance dramatically. (J. Wang et al. 2016) 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). Those previous research were based either on qualitative and regulatory analysis (Lansford, Kenney, and Ecclesine 2013; K.-H. Chang 2015), analytical methods (J. Wang et al. 2016; Liu, Naik, and Park 2017), and/or simulation (Liu, Naik, and Park 2017; Cheng et al. 2017; Park and Kim 2014). One exception also employed a limited experimental setup (Naik, Liu, and Park 2017).

The scenarios and assumptions adopted in previous work are very different from ours (our scenario and assumptions are described in Section 5.3). Nevertheless, their conclusions are consistent with our findings described in Section 5.4, regarding how each type of device affects the other performance. However, our work is novel in the sense that we combine an analysis of sharing with both an assessment of how much spectrum to allocate, and whether cooperation between V2X and unlicensed devices is effective.

5.3. *Method and system model*

We employ the engineering-economic approach described in Chapter 2, which has been modified with the extensions necessary to address the spectrum-related issues presented in this chapter. One major part of the analysis is to use packet-level simulation to examine how the ability of vehicular networks and hotspots to carry IP traffic is affected by sharing spectrum between those types of devices. For such an analysis, we adopt the simple measures of performance of throughput to vehicles and throughput to unlicensed devices connected to hotspots, for cases where data rates of incoming traffic, i.e. the total data rates demanded by the devices, are fixed. For this simulation, data from Wi-Fi hotspots operating in Portugal is used together from the data from the vehicular network described in Chapter 2, to define the simulation parameters. Several factors are varied in the simulation to observe the effect of each on throughputs to V2X and unlicensed devices. One factor is the amount of spectrum used, for which we either assume vehicles and hotspots use separate spectrum, or we assume spectrum is shared between vehicles and hotspots. Sharing can take place according with different sharing schemes, with devices coexisting in a co-equal basis, using 802.11 listen-before-talk mechanisms to mitigate mutual interference, or with unlicensed cooperating with V2X devices to relay traffic from the latter. The throughput simulation lets us address both the issues of how much spectrum to allocate for ITS, whether that spectrum should be shared with unlicensed devices, and if so, under what sharing scheme. Other factors that are varied in the simulation include the densities of vehicles and hotspots, data rates of incoming Internet traffic to vehicles and hotspots, and whether hotspots are located indoors or outdoors. Another part of our method is to use the vehicular simulated throughput to estimate the economic benefit of adding ITS spectrum to offload Internet traffic and assess whether that benefit exceeds the opportunity cost of ITS spectrum. The model, data used, and assumptions are described below.

This remaining of this section is organized as follows. First, we describe the model of usage and sharing of the ITS spectrum, as well as the economic model and its assumptions. The simulation model with the assumptions for the vehicular and hotspot throughputs are then described, as well as the Portugal dataset used to set simulation parameters. Finally, we describe the base-case assumptions that apply for the results presented in this chapter.

5.3.1. 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 by those types of devices and what benefits are considered.

Like the coexistent sharing-among-equals proposal to the FCC, 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. This is consistent with (Qualcomm 2013), which proposes that three 10 MHz channels in the ITS band be allocated exclusively for safety. Our model allows us to evaluate non-safety benefits from adding spectrum, in a way that is independent from whatever safety benefits are achieved. In the 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 to carry IP traffic as follows. For the vehicular network, bidirectional connections are established between each vehicle equipped with a V2X onboard unit (OBU) and one fixed RSU which serves as a gateway to the Internet, as described in more detail in Chapter 2. Each vehicle uses one channel chosen from a number $D \geq 0$ of channels, while each RSU can use all D channels. Each hotspot uses one channel chosen 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, we assume that V2X and unlicensed devices share spectrum through the listen-before-talk mechanisms specified in IEEE 802.11, the standard used by both DSRC and unlicensed devices such as Wi-Fi. We assume sharing on a co-equal basis, i.e. V2X and unlicensed devices have equal priority when transmitting. V2X and unlicensed devices detect each other's transmissions by means of preamble detection, which means that devices are able to decode the packets of the other type of device. This is the type of transmission detection considered in one of the proposals to the FCC, as opposed to energy detection[12].

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 (Peha 2009) and how. These schemes are illustrated in Figure 5.1. The simplest of the three schemes is what we call *coexistence* without cooperation. With this scheme, V2X and unlicensed devices sense each other's transmissions, but devices of one type try to avoid interference without explicit cooperating with devices of the other type (Peha 2009).

In contrast, the two other schemes require that hotspots cooperate to relay V2X packets, in order to be allowed to operate in shared spectrum. In one of these cooperation schemes, hotspots act as V2X RSUs. Thus, vehicles that are not in communications range with RSUs or other vehicles can send packets to hotspots, which relay those packets to and from the

12 Preamble detection means that devices are able to detect the bits used to set and synchronize the receiver in order to receive packets. This contrasts energy detection, in devices sense the power of transmissions from others, but are not able to decode the transmitted packets. In IEEE 802.11 standards, preamble detection should work for received signals at lower power than with energy detection, meaning that preamble detection is able to detect nodes at relatively farther distances than with energy detection.

backhaul connection to the Internet attached to those hotspots. We call this scheme *backhaul cooperation*.

In the second cooperation scheme, hotspots act like V2X devices in vehicles. This scheme differs from the one above in the sense that hotspots do not relay V2X traffic into their backhaul connections. Rather, hotspots relay traffic between V2X devices within the ITS band. This is also important for a vehicle that is not in range with RSUs or other vehicles but can use hotspots as non-moving relays in a multihop route to reach an RSU. We refer to this strategy as *relay cooperation*.

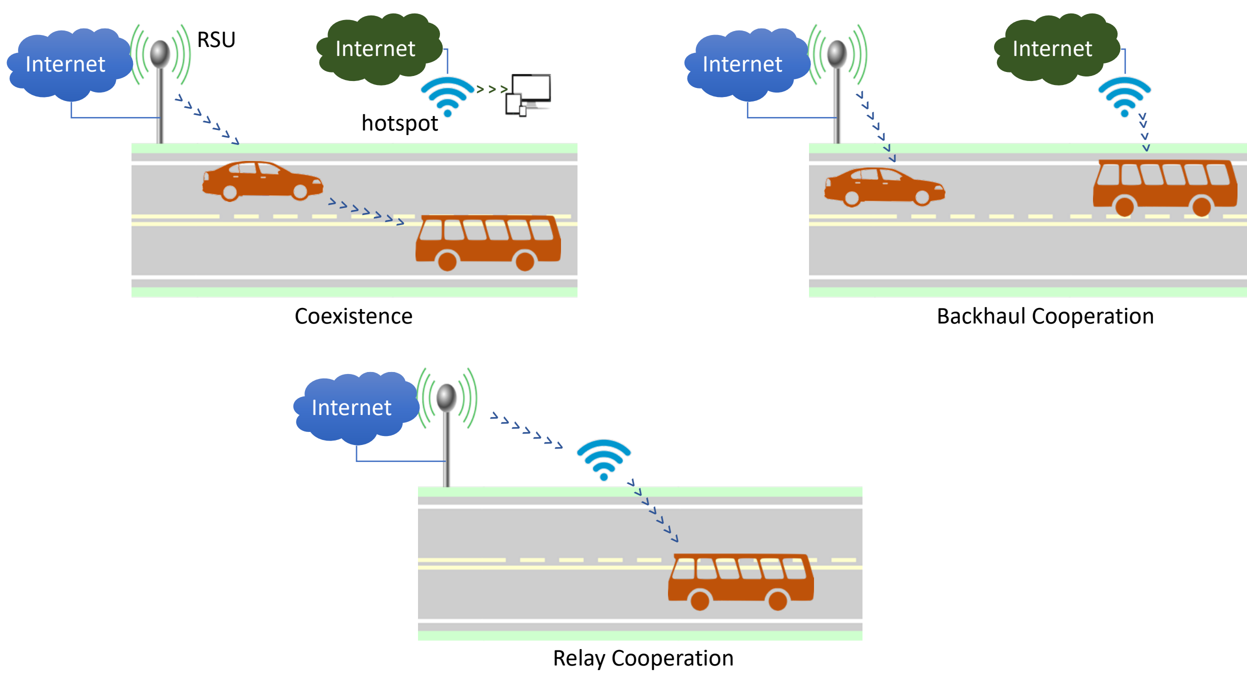

Figure 5.1. Representation of coexistence, backhaul cooperation and relay cooperation sharing schemes.

We assume that each channel is 10 MHz wide, which is the current channel specification in the ITS band. For all sharing schemes, 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. The assignment of channels to vehicles and hotspots is as follows. First, one channel is assigned at random per hotspot, with all channels having equal assignment probability. Then, channels are assigned to vehicles, one at a time. If hotspots are operating under one of the cooperation sharing schemes, and if a given vehicle is connected to a hotspot, then that vehicle is assigned to the same channel assigned to the hotspot it is connected with. Otherwise, the channel to be used by each vehicle is chosen according to the method proposed in (Ramachandran et al. 2006). That method takes into consideration the expected interference from nodes already assigned to the channels. When a channel is used by either V2X devices only or unlicensed devices only, all interferers are of the same type. When otherwise a channel is shared, interferers can be both V2X and unlicensed devices.

In the channel assignment model above, a multihop route is established only if all nodes (vehicles and hotspots, in case of cooperation) are assigned to the same channel.

5.3.2. 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 consider that V2X networks have the ability to offload mobile Internet traffic otherwise carried by capacity-limited microcellular networks. In Chapter 2, we define the benefit of offload as the cost savings from deploying fewer cell towers. Social welfare is maximized when the marginal benefit per unit of spectrum added equals the marginal costs of offloading (Hazlett and Honig 2016). 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.

Costs can be of three types. Two are the costs of in-vehicle devices and RSU infrastructure, which are described Chapter 2. Regarding the latter, our simulations suggest that in a given scenario the quantity of RSUs that maximizes benefit minus cost is approximately insensitive to bandwidth. Hence, we keep that quantity fixed when spectrum amount is varied in a given scenario. The other cost 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 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 (or MHz-pop) for bands in 1.8-2.2 GHz in 2015, which were considerably more expensive than bids that paid around $0.60 for similar frequencies in 2006 (Aittokallio 2015). It must be taken into consideration that physical properties of spectrum make it far less valuable at higher frequencies (e.g. 5.9 GHz) than at lower frequencies (Peha 2013b; Alotaibi, Peha, and Sirbu 2015), perhaps by an order of magnitude. That might place the value of ITS spectrum in the order of a few tens of cents. However, emerging technology operates effectively at higher technology than was typical in the past, so the value of higher frequencies is probably changing, which adds to its uncertainty.

Moreover, current use of Wi-Fi at 5 GHz and the FCC NPRM on sharing indicate that ITS spectrum might be opened for unlicensed use. Estimating the marginal value of unlicensed spectrum is very difficult, but marginal value per MHz would certainly be less than value per MHz averaged over all spectrum. A group of organizations interested in expanding the use of unlicensed spectrum has estimated the total value of spectrum (Katz 2014), which 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.

5.3.3. Simulation model and assumptions

Our method depends on estimates of throughputs to address the research questions presented in this chapter. We simulate throughputs at packet-level from the physical to the transport layer, and the part of the simulation model that represents vehicles and RSUs is described in greater detail in Chapter 2. We have extended the model to vary the amount of ITS spectrum and to allow sharing with unlicensed devices. A different simulation is run for every scenario of numerical assumptions and spectrum strategy (separate bands, coexistence, backhaul cooperation or relay cooperation).

Unlicensed devices are represented in the simulation by Wi-Fi hotspots and Wi-Fi devices, which should be the majority, if not all, of the unlicensed devices that will likely share the ITS band. 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 hotspots, such as those in residences and offices, and outdoors hotspots, such as those for public Wi-Fi in open locations. For indoor hotspots, we assume that in any given 5-s interval some hotspots are active while others are not. Active hotspots are receiving packets at a constant rate throughout the 5 seconds, while inactive hotspots receive no packets. Every 5 seconds, a different set of hotspots is randomly selected to be active. Moreover, we assume that 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, Portugal (see subsection below about the dataset). 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 and shifted 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. This way we obtain samples with desired hotspot density, and with coordinate distribution which intensity approximates that of the original set. We assume 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 is assumed to propagate according to an indoor propagation model (Meinilä et al. 2009) to the endpoint of its TCP connection, or a model with wall obstruction with V2X devices or outdoor hotspots.

The assumptions for outdoor hotspots are different. These are placed along the streets of Porto (see subsection below about the dataset). In a given street, the inter-hotspot distance is fixed. Signal propagates according to the same outdoors loss model used for vehicles and RSUs. Moreover, we assume that all outdoor hotspots are active at peak hours. The transmission power of all hotspots and their clients is 11 dBm at the antenna output, which is consistent with popular Cisco Wi-Fi hotspots (Cisco 2014).

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

5.3.4. *Portugal dataset*

To set some of the simulation parameters, we use data from a real vehicular network operating in Porto, Portugal, as described in Chapter 2. To address the spectrum-related research questions presented in this chapter, we also use data from Wi-Fi hotspots and the coordinates of roads in that city. We have collected positions of 65,000+ Wi-Fi hotspots in Porto, which were available in the website of FON, one major Wi-Fi service provider. The dataset includes Wi-Fi hotspots from the subscribers of a major fixed broadband provider in Portugal who partners with FON. 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. Second, strength of the signal received from RSUs is verified to be compatible with the simulated signal strength in vehicles and RSUs, 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 the previous subsection.

5.3.5. *Base-case numerical assumptions*

The results of the analysis presented in this chapter refer to a base-case scenario of numerical assumptions, of scenarios derived by varying one or more assumptions from the base-case scenario. Some of those base-case numerical assumptions are those presented in Chapter 2. The base-case assumptions that are specific to address the research question presented in this chapter are described as follows.

The base-case assumption for the penetration of V2X OBUs in vehicles of 100%. This is reasonable over the timeframe of a spectrum allocation decision if the Department of Transportation mandates V2X for safety communications (U.S. Department of Transportation 2016). We also examine the impact of lower penetrations on our results, as might be appropriate if no mandate occurs.

Another assumption that is highly uncertain is the data rate per vehicle. We assume a "low" case value of 400 kbps that is consistent with (Ligo et al. 2017), but we also present results for much higher data rates, because data rates have been increasing rapidly over time (Cisco 2016), and future data rates are uncertain.

For other values, we use base assumptions that are representative of five years into the future. Although this work informs spectrum allocation decisions that may span decades, the rate of technological change and adoption in wireless communications make decade-long predictions highly uncertain. Since five years is a typical horizon for predictions about Internet usage for given technologies – see e.g. (Cisco 2016) –, we adopted five years as our horizon for analysis.

The base assumption for the average data rate of incoming Internet traffic in the peak hour over active hotspots is 5 Mbps in five years. This value is reasonable because it has been found that the majority of traffic in the U.S. is currently from video applications (Engebretson 2016). Typical video streams have an average bitrate of 2 Mbps – see (Ozer 2016), Netflix HD encoding –, which we assume as today's average peak-hour data rate per active hotspot, and usage for fixed broadband subscribers is forecast to grow at roughly 19% per year (Cisco 2016).

We assume that the quantify of indoor hotspots in a location is proportional to population density, at a rate of one hotspot for every four people, which is based on (CIsco 2017). (Also, see subsection 5.3.3 for the method to place indoor hotspots.) Moreover, we assume that 15% of indoor hotspots will be active at a time. This is reasonable because current estimates for the average traffic in U.S. households are currently around 100 GB per month (AT&T 2017), (Cisco 2016), (Engebretson 2016). A hotspot transferring 300 kbps at all times would transfer 100 GB over a month, then the share of active hotspots is assumed as 300 kbps / 2 Mbps = 15%. Although this assumption about the share of active hotspots at a given time several years 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 results, 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 downtown areas of a few cities and other sparse locations, this assumption is also likely to result in higher interference to vehicles than in typical urban areas. (For this reason, in the results section we compare scenarios with both indoor and outdoor hotspots with scenarios with indoor hotspots only.)

For the results that refer to a specific location, the base-case population density is 2,000 people per km^2, which is representative of a city like Pittsburgh, unless stated otherwise.

5.4. Results and discussion

In this section, we first address the issue of how much spectrum to allocate for ITS. We then address the issue of whether the ITS band should be shared with Wi-Fi devices, and if so, under what rules (coexistence only or with cooperation as defined in subsection 5.3.1). The throughput for each scenario of bandwidth, device density and data rates is derived by averaging throughput for more than 1,000 vehicles and a larger number of hotspots. Assuming that the throughputs are mutually independent, then the 95% confidence interval is within 5% of the mean throughput for vehicles and 2% for hotspots.

5.4.1. How much spectrum to allocate for ITS

To address the issue of how much spectrum to make available, 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.

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. for those locations 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. the optimal quantity of RSUs to deploy (or not) for Internet access is determined at each census tract based on its average population density – this approach was also employed in (Hallahan and Peha 2011b, 2010). (The optimal quantity of RSUs in each location is defined as the quantity that maximizes benefit minus RSU cost in that location. See Chapter 2.)

Figure 5.2 shows marginal and average benefit minus RSU cost (B-C) per MHz-pop on a nationwide scale for the U.S., as a function of bandwidth allocated exclusively for vehicles. The graph shows results for two data rates of incoming Internet traffic per vehicle (low and high scenario as defined in subsection 4.3.5). The other assumptions are base-case values. For a particular bandwidth to be worth allocating, both marginal and average benefit minus RSU cost B-C must exceed the opportunity cost of ITS spectrum. If marginal benefit minus cost is less than opportunity cost at bandwidths where the former is decreasing, then reducing bandwidth increases benefit minus cost. If average benefit minus cost is less than opportunity cost, then benefit minus cost is greater with a bandwidth of 0.

Figure 5.2 shows that benefit minus RSU cost 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 deployed in an area (see the expression for RSU cost in Chapter 2). We found that the optimal quantity of RSUs deployed is roughly invariant with bandwidth, for the range of data rates, densities and other factors we considered.

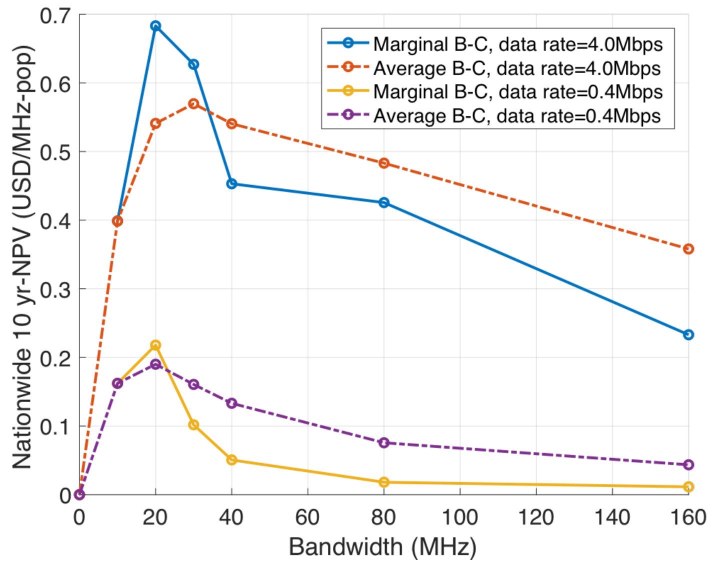

Figure 5.2. Nationwide benefit minus RSU cost per capita (B-C), as a function of bandwidth. Lines for two different data rates of incoming traffic per vehicle are shown. OBU penetration and other assumptions are at base-case values.

Given the uncertainty in the opportunity cost of ITS spectrum as discussed in subsection 5.3.2, we examine the relationship between the opportunity cost and the optimal bandwidth in Figure 5.3. For a given opportunity cost, the graph shows the maximum bandwidth for which marginal and average benefit minus RSU cost exceed that opportunity cost. In subsection 5.3.2 we conjecture that the cost of ITS spectrum might be around \$0.20-\$0.40 per MHz-pop. Figure 5.3 shows that for such a range of opportunity cost it might be worth allocating spectrum, but the amount that maximizes social welfare depends not only on spectrum cost but also on other factors as well, such as data rates. For example, at an OBU penetration of 100% of vehicles and average data rate of incoming traffic of 4 Mbps per vehicle, Figure 5.3 shows that it is worth allocating 40 MHz of ITS spectrum, which is the bandwidth currently available for non-safety use, 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 the opportunity cost is much lower (below $0.05 per MHz-pop).

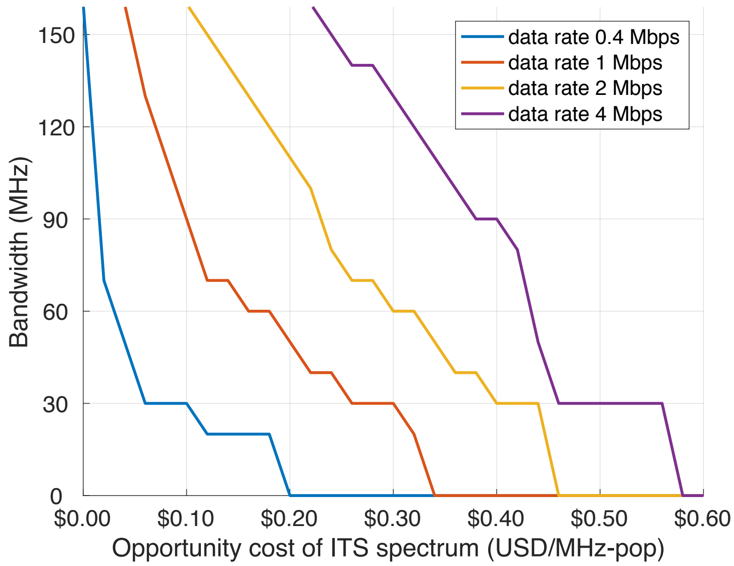

Figure 5.3. Bandwidth that maximizes social welfare (on a nationwide basis), as a function of the opportunity cost of spectrum in the ITS band. Curves are shown for distinct data rates of incoming traffic per vehicle. The other numerical assumptions are at base-case values.

Moreover, the results above are for an OBU penetration of 100%, which is consistent with a mandate of V2X in all vehicles. Out of the context of a mandate, lower penetrations are possible, with OBUs being deployed more frequently in vehicles that demand higher data rates. Figure 5.4 shows the bandwidth that maximizes social welfare in such a scenario. The graph 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 Figure 5.4 than for the scenarios with 100% penetration (Figure 5.3). However, Figure 5.4 shows that a small increase in penetration (5% to 10% in the graph) changes significantly the bandwidths worth allocating, especially if the opportunity cost of spectrum is in the order of tens of cents per MHz-pop (as discussed in subsection 5.3.2). For example, at an OBU penetration of 10% it is worth allocating 40 MHz (the bandwidth currently available for non-safety use) if the cost of spectrum is about \$0.18 per MHz-pop. However, a scenario where it is not worth allocating spectrum in excess of safety is also plausible, especially for low OBU penetrations and/or if spectrum is valued at more than a few tens of cents per MHz-pop.

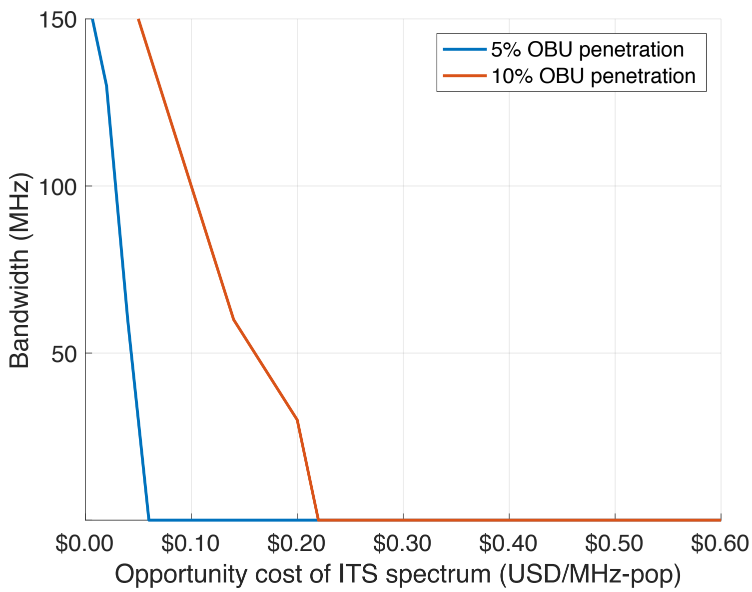

Figure 5.4. Bandwidth that maximizes social welfare (on a nationwide basis) as a function of the opportunity cost of spectrum in the ITS band. Curves are shown for distinct penetrations of OBUs. Data rate per OBU is 27 Mbps (the maximum for 802.11p in a 10 MHz channel) and other assumptions are at base-case values.

If more spectrum allocated for ITS means less data traffic being carried over macrocells, then less towers are needed in the cellular network. The benefit of less towers is modeled in Chapter 2 – 2.3, and depends on factors such as cellular bandwidth, spectrum efficiency and cost per tower. Figure 5.5 shows the bandwidth that maximizes social welfare for the base-case tower cost ($750,000) and a lower cost. The graph shows that the amount to allocate for a given opportunity cost of spectrum varies greatly with tower cost. For example, if the cost of spectrum is around $0.20-0.40 (as discussed in subsection 5.3.2) and data rates are high, then it is worth allocating 90+ MHz for the base-case value of the average tower cost. However, it is not worth adding any spectrum for offload if tower cost is about 30% lower ($550 thousand in Figure 5.5). The same applies if macrocell bandwidth or spectrum efficiency is 30% more than the base values presented in Chapter 2 – 2.4.

In this subsection we see that there are realistic scenarios in which it is worth allocating more spectrum than it is currently available for non-safety (40 MHz), but there are also scenarios in which it is not worth adding any spectrum. However, it is important to note that this discussion applies for spectrum allocated exclusively for ITS. While we found that the bandwidth that maximizes social welfare depends on the uncertainty of its opportunity cost and other factors, our estimates of benefit do not capture the value of sharing spectrum with unlicensed devices. Benefits of sharing are discussed in the following subsection.

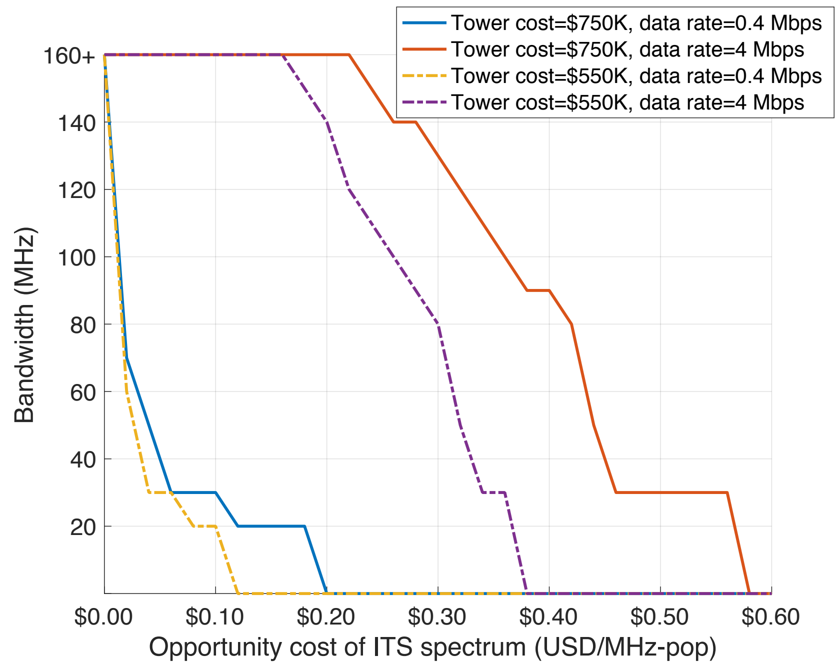

Figure 5.5. Bandwidth that maximizes social welfare (on a nationwide basis) as a function of the opportunity cost of spectrum in the ITS band. Curves are shown for distinct data rates and tower cost. Other numerical assumptions are in base-case values.

5.4.2. Whether to share ITS spectrum with unlicensed devices and how

In this subsection we address the issue of whether and how to share spectrum. We do this by comparing performance of four different strategies: allowing vehicles and unlicensed devices to share spectrum with the three possible sharing schemes defined in subsection 5.3.1, and placing vehicles and unlicensed devices in separate bands. That comparison is done with respect to two measurements. The first is a comparison among the strategies with respect to throughputs to vehicles and unlicensed devices. The second is a comparison among the strategies with respect to the amount of spectrum needed to achieve given throughputs.

5.4.2.1. Throughput to vehicles and unlicensed devices

Figures 5.6 and 5.7 show vehicle throughput and hotspot throughput for the four different strategies. In both graphs the horizontal axis is the bandwidth allocated (in excess of what is used for safety). We show throughputs for indoor hotspots only, and for both indoor and outdoor hotpots. As expected, throughput increases with spectrum bandwidth for both vehicles and hotspots, although at a diminishing rate.

The graphs also show how throughputs differ among strategies. Figure 5.6 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 Figure 5.6 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 (Figure 5.6, right). Because of the interference from outdoor hotspots, for any bandwidth, throughput to vehicles coexisting with unlicensed devices is significantly lower than throughput without sharing.

The loss of vehicle throughput caused by sharing can be mitigated if unlicensed devices are required to cooperate, as shown in Figure 5.6 (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 increases. However, not all cooperation schemes increase vehicle throughput (relative to coexistence). Figure 5.6 shows that backhaul cooperation, in which hotspots act as V2X RSUs by relaying V2X packets through hotspot backhaul, increases vehicle throughput when there are outdoor hotspots (right). However, relay cooperation, which implies that hotspots relay packets between vehicles and RSUs, does not increase throughput (The resulting throughput from the additional connections, if any, does not exceed the additional interference from those connections with relay cooperation).

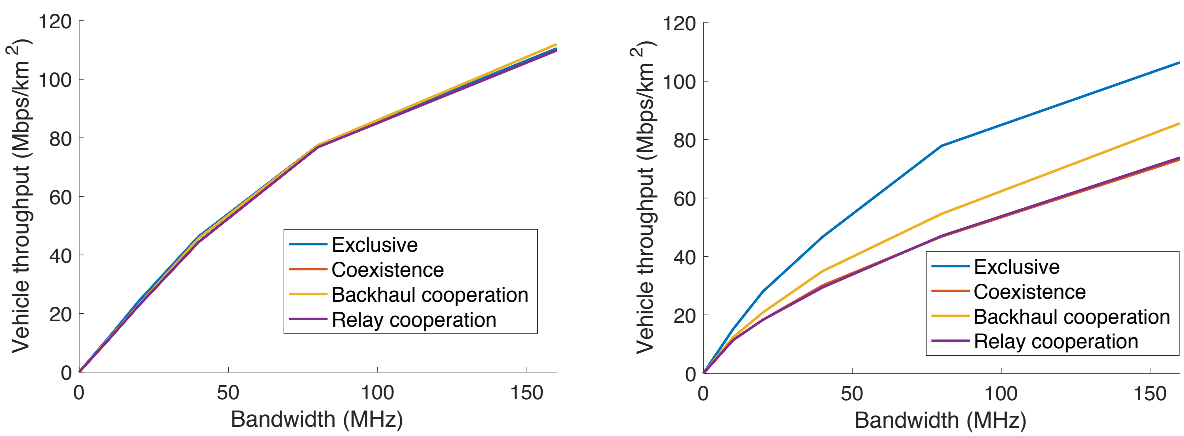

Figure 5.6. Vehicle throughput under different sharing strategies. The left graph refers to indoor hotspots only, the right graph is for indoor and outdoor hotpots. The other assumptions are at base-case values (with "high" data rates).

On the other hand, hotspot throughput is relatively less sensitive to sharing (and the sharing scheme) than vehicle throughput. Figure 5.7 (left) shows that for indoor hotspots the difference in hotspot throughput among all strategies is negligible. (All curves overlap.) Even when outdoor hotspots are present (right graph in Figure 5.7), the difference between hotspot throughput in exclusive spectrum and in shared spectrum is small, and the curves for the sharing schemes still overlap. The impact of sharing on hotspot throughput is probably small because the densities and data rates of hotspots are much higher than that of vehicles for the scenario shown, and many of the hotspots are indoors, which are separated from streets by walls.

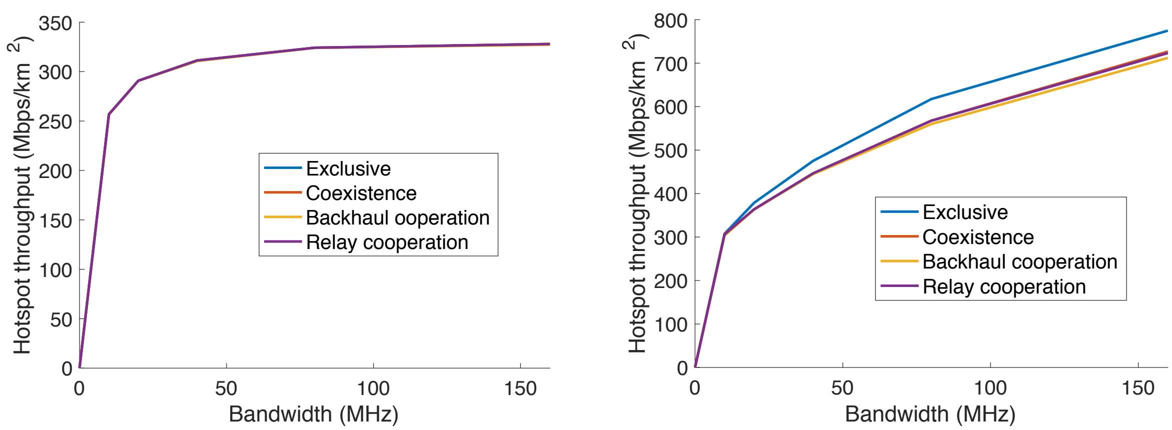

Figure 5.7. Hotspot throughput under different strategies. The left graph refers to indoor hotspots only, the right graph is for indoor and outdoor hotpots. The other assumptions are at base-case values (with "high" data rates).

Figure 5.6 indicates that the presence of outdoor hotspots has significant impact on vehicle throughput. Hence, in Figure 5.8 we examine throughputs for varying densities of outdoor hotspots. The horizontal axis shows increasing density, i.e. decreasing distance between outdoor hotspots (the base-case value is of 150 m between hotspots). The left graph shows that vehicle throughput with a higher separation between outdoor hotspots is significantly less than throughput with a lower separation (i.e. higher density). For example, when there are as three times as more outdoor hotspots as in the base case, vehicle throughput with coexistence is 2/3 lower than on exclusive spectrum.

The impact of outdoor hotspots on vehicle throughput can be mitigated with backhaul cooperation. Fig. 7 (left) shows that throughput is higher with backhaul cooperation than with coexistence for all densities examined. For lower hotspot densities (300 m separation or more), vehicle throughput with backhaul cooperation is almost the same as with exclusive spectrum. However, this throughput “advantage” of backhaul cooperation over coexistence is smaller with higher densities of outdoor hotspots. This is because the benefit in backhaul cooperation of relaying vehicular traffic does not change with density, whereas interference from hotspots to vehicle increases.

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. Increasing hotspot density results in more interference to vehicles for sharing schemes. However, relay cooperation does not result in more throughput from hotspots relaying V2X packets between RSUs and other vehicles.

Figure 5.8 (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 no significant difference in the burden that these strategies impose on hotspot throughput.

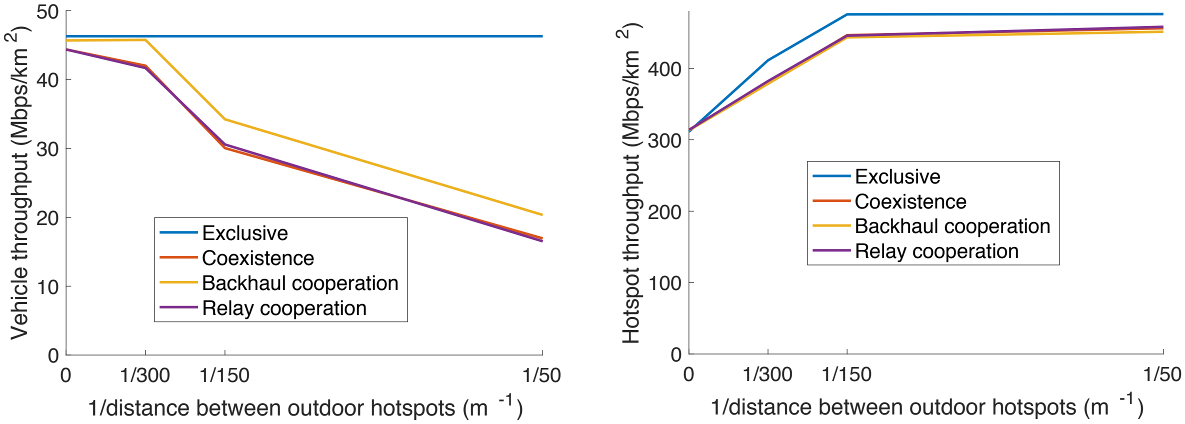

Figure 5.8: Throughputs to vehicles (left) and hotspots (right) as a function of decreasing distance between outdoor hotspots, under different strategies. Bandwidth is 40 MHz, and the other assumptions are at base-case values (with "high" data rates).

In Figures 5.6 to 5.8, we examine how throughput varies with density of outdoor hotspots, for different sharing strategies. In the following we also analyze how those throughputs vary with other factors, such as data rates and densities of devices other than outdoor hotspots. For all those scenarios, we find that the differences between the throughputs on shared spectrum and the throughputs on separate spectrum are not as large as the scenarios of Figure 5.8. Moreover, the differences in throughputs among the different sharing strategies are similar to those found in Figure 5.8.

Figure 5.9 shows throughputs for varying data rates of vehicular Internet traffic. (The "low" base-case data rate is of 0.4 Mbps/vehicle while the "high" base rate is 4 Mbps/vehicle). The left graph shows that for the "low" data rate, vehicle throughput with coexistence is the same as throughput without sharing. This is because at this data rate vehicles face less congestion in the channels than at higher data rates, even with hotspot interference. Nevertheless, backhaul cooperation achieves higher throughput than relay cooperation and coexistence, because there are vehicles communicating with the Internet through backhaul cooperation that were disconnected from RSUs in the other strategies. Because of this, backhaul cooperation achieves even higher throughput than spectrum that is not shared with hotspots. The same does not occur with relay cooperation, however. The presence of hotspots relaying traffic between other vehicles and RSUs does not increase throughput.

As data rate increases, Figure 5.9 (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 with unlicensed devices than on spectrum allocated exclusively. Still, backhaul cooperation results higher vehicle throughput than coexistence, while relay cooperation results roughly the same throughput as coexistence. (The differences between the coexistence and relay cooperation curves are within the confidence intervals of the simulations.)

Figure 5.9 (right) shows that throughput to unlicensed devices with higher vehicle data rates is less than throughput with lower data rates, which is expected 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, Figure 5.9 (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.

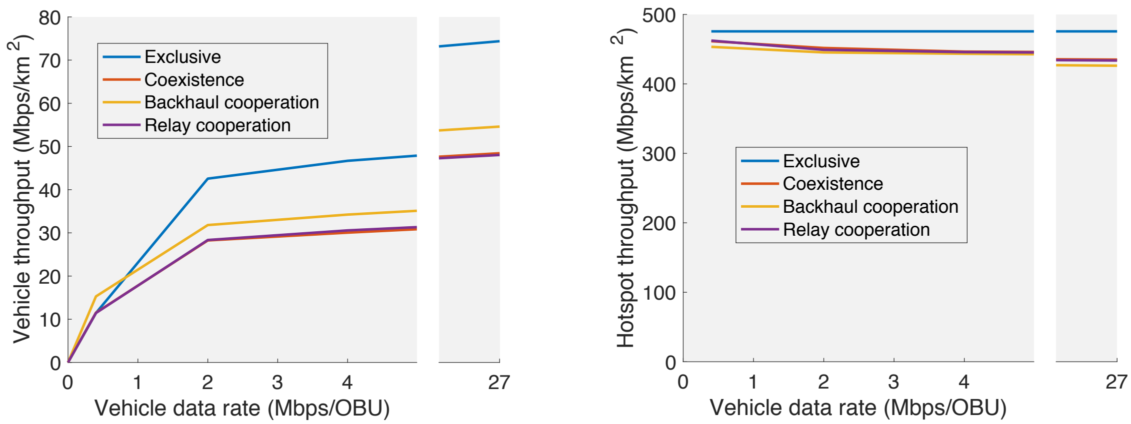

Figure 5.9. Throughputs to vehicles (left) and hotspots (right) as a function of data rates of Internet traffic per vehicle, with no sharing and with different strategies. Bandwidth is 40 MHz, and the other assumptions are on base-case values.

Data rates are expected to increase over time. Likewise, the penetration of V2X devices in vehicles is also expected to increase, resulting in higher densities of those devices over time. Figure 5.10 shows that the effect of increasing V2X penetration on throughputs for different sharing strategies is similar to those previously shown for increasing data rates. The left graph shows that for any vehicle density, vehicle throughput with any sharing strategy is significantly less than throughput without sharing, while backhaul cooperation results in higher vehicle throughput than coexistence, and relay cooperation results in roughly the same throughput as coexistence.

Figure 5.10 (right) shows that throughput to unlicensed devices decreases with V2X penetration, because of the increasing quantity of vehicles, although that decrease is relatively small. The graph also shows that hotspot throughput with all sharing schemes is slightly less than throughput with exclusive spectrum, and the hotspot throughput is not significantly different among sharing schemes.

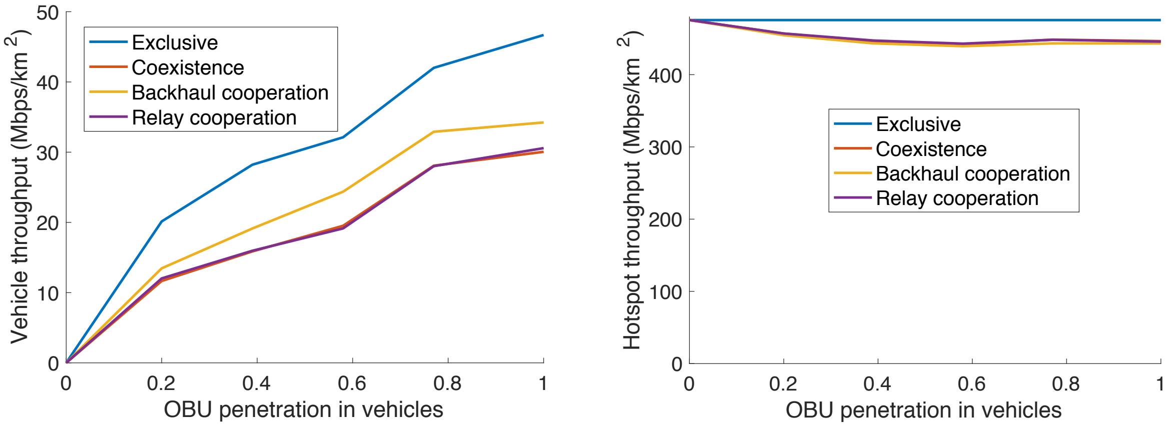

Figure 5.10. Throughputs to vehicles (left) and hotspots (right) as a function of V2X penetration, with no sharing and with different strategies. Bandwidth is 40 MHz, and the other assumptions are at base-case values (with "high" data rates).

Although Figures 5.6 to 5.10 showed that backhaul cooperation results in higher vehicle throughput than with coexistence, for a wide span of conditions of data rates and device densities, the difference between the sharing strategies may disappear for extremely high densities of devices. Figure 5.11 (left) shows the effect of increasing population density (which results in higher densities of both vehicles and hotspots) on vehicle throughput for different sharing strategies. The graph shows that vehicle throughput with any sharing scheme is less with higher population densities, because of the higher number of interfering devices of both types. Moreover, backhaul cooperation results in higher vehicle throughput than with coexistence for most population densities, but the difference diminishes for densities around 7,000 people/km^2 or more. This is 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, it is also true that this density is extremely high. (Only a very few locations, e.g. Manhattan, have densities higher than 7,000 people/km^2.) Therefore, the differences between the sharing strategies shown for densities below 5,000 people/km^2 should hold for the vast majority of the U.S.

Figure 5.11 (right) shows that for population densities around 1,000 people/km^2 or less, there is no significant difference between hotspot throughput with any sharing strategy 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 all sharing strategies is slightly less than throughput with exclusive spectrum (although the difference is statistically significant). However, this difference is negligible for densities around 7,000 people/km^2 or more. This is again because the quantity of unlicensed devices increases with population density faster than the quantity of vehicles. Hotspot throughput with any strategy is not significantly different from the other sharing strategies.

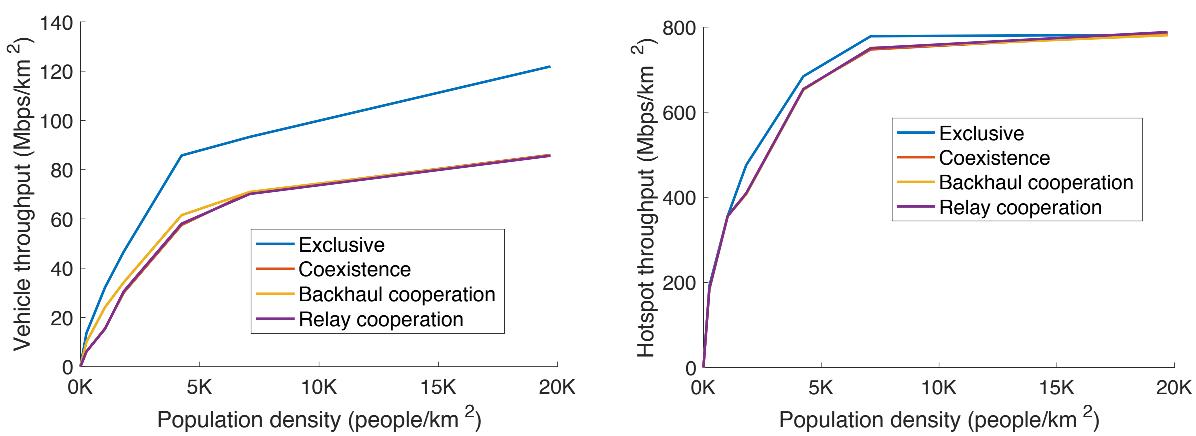

Figures 5.11. Throughputs to vehicles as a function of population density, with no sharing and with different strategies. Bandwidth is 40 MHz, and the other assumptions are on base-case values (with "high" data rates).

For a sharing scheme to be the best, it must achieve throughputs for the two device types that other sharing schemes cannot. Figure 5.12 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 between zero and the total amount. 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* (Peha and Tobagi 1996), 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 spectrum-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 be achieved with backhaul cooperation but not with the other sharing schemes. That is, for the numerical assumptions used for Figure 5.12, vehicle throughput higher than 17 Mbps/km^2 can only be achieved with backhaul cooperation, or by increasing the total amount of spectrum used.

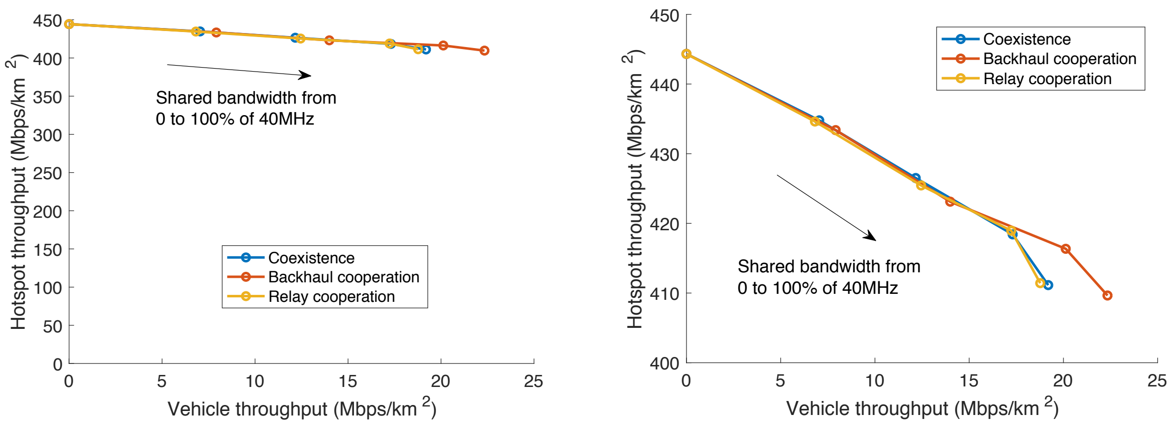

Figure 5.12. Hotspot throughput as a function of vehicle throughput, with different sharing schemes. Bandwidth is fixed at 40 MHz, and the other assumptions are at base-case values (with "high" data rates).

The results above are related to previous research described in Section 5.2, because we investigate whether to share, which includes an analysis of throughput performance with and without sharing. However, the research questions, scenarios addressed, and results presented in this section are different from previous work. First, the assumptions and results work such as (Naik, Liu, and Park 2018, 2017; J. Wang et al. 2016; Cheng et al. 2017; Liu, Naik, and Park 2017; Park and Kim 2014; Lansford, Kenney, and Ecclesine 2013; K.-H. Chang 2015) are more applicable to DSRC for safety communications, while our work applies to Internet traffic. This is because they consider DSRC devices broadcasting data, which is typical for safety applications. In contrast, we consider unicast connections over a mesh network formed among vehicles and Internet-connected RSUs (sometimes cooperating with hotspots, depending on the sharing scheme), over which TCP connections are established to carry Internet traffic. Second, previous work such as (J. Wang et al. 2016) 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 (J. Wang et al. 2016) 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.

Despite the differences between previous work and ours, most of the papers above found that Wi-Fi devices can interfere significantly with the performance of DSRC devices, especially if the former operate in saturated mode, i.e. when every Wi-Fi device always have a packet to send, or when Wi-Fi devices aggregate multiple adjacent channels. However, most of those authors also concluded that if the interval between Wi-Fi packets is above a certain threshold, performance of DSRC devices approach that in the absence of sharing. Although the scenarios and assumptions are very different from ours as described above, their conclusions are consistent with our findings regarding how each type of device affects the other performance.

5.4.2.2. Required bandwidth to achieve given throughputs

In this section, we determine 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 the vehicle throughput and the hotspot throughput 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 vehicular and hotspot throughputs.

Figure 5.13 shows the amounts of spectrum obtained with the procedure above as a function of vehicle and hotspot throughputs. One curve refers to the total amount of spectrum when vehicles and hotspots use spectrum separately, and the others to the amount of spectrum with different sharing schemes. (The curves for coexistence and relay cooperation overlap.) The horizontal axis represents vehicle throughput and the colors represent hotspot throughput. The curves for a given vehicle throughput also refer to the same hotspot throughput (i.e. the curves at any given vehicle throughput have the same color).

For the base-case scenario, Figure 5.13 shows that 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, the graph shows that 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 than the other schemes 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.

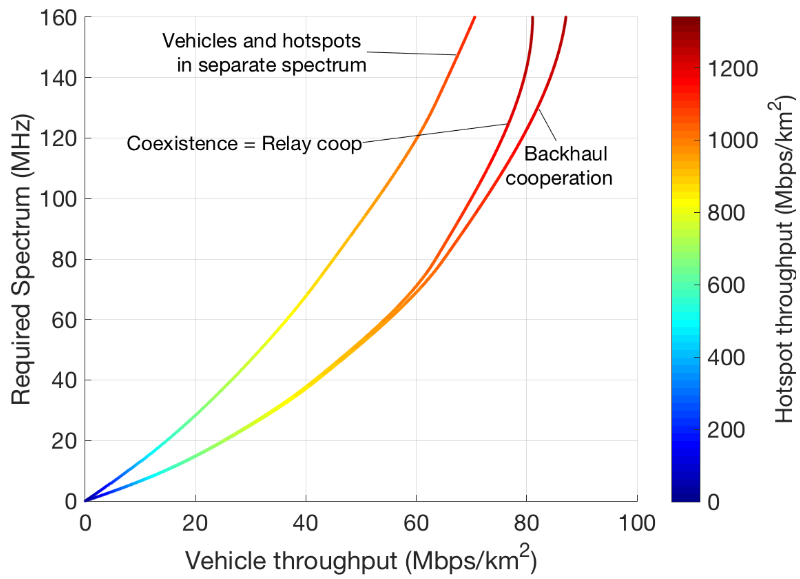

Figure 5.13. Required spectrum to achieve given vehicular and hotspot throughputs, as a function of vehicular throughput, for different strategies. Points of equal color refer to equal hotspot throughput. Colors are coded in the bar (right). Assumptions are at base-case values (with "high" data rates).

We find that less spectrum is required when that spectrum is shared for other scenarios as well. In particular, Figure 5.14 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. Figure 5.14 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 Figure 5.14 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.

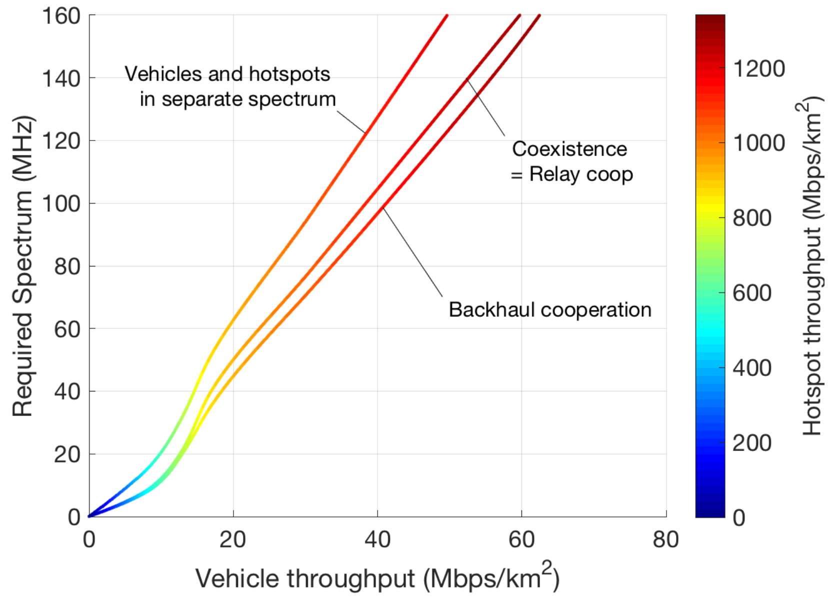

Figure 5.14. Required spectrum to achieve given vehicular and hotspot throughputs, as a function of vehicular throughput, for different strategies. Distance between outdoor hotspots is 50 m, and the other assumptions are at base-case values (with "high" data rates).

The effect of population density on required bandwidth for different sharing strategies is not obvious, because quantities of both V2X and unlicensed devices vary with population density. For this reason, we examined the required bandwidth for lower and higher population densities than previously shown. Figure 5.15 shows the required bandwidth to achieve given throughputs for a location with 250 people/km^2. For this population density, we find that 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. The left graph shows that coexistence requires significant less bandwidth than vehicles and unlicensed devices using separate bands, when there are no outdoor hotspots. This is because there are less devices than in the previous scenarios, and indoor hotspots do not interfere significantly with vehicles. For the same reason, cooperation has little effect in that scenario. On the other hand, Figure 5.15 (right) shows that the presence of outdoor hotspots results in almost no bandwidth savings of coexistence when compared to vehicles and unlicensed devices using separate bands. This is because of the significant interference caused by hotspots on vehicles. However, backhaul cooperation results significant bandwidth savings, because the interference is partially mitigated by extra vehicles being connected to the Internet through hotspots.

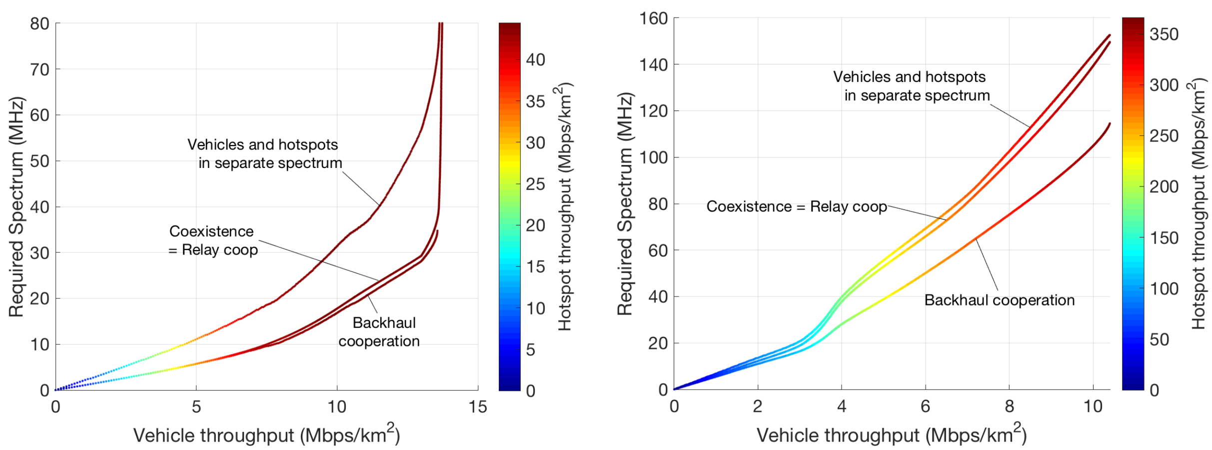

Figure 5.15. Required spectrum to achieve given vehicular and hotspot throughputs, as a function of vehicular throughput, for different strategies. The left graph is for indoor hotspots only, and the right graph is for indoor and outdoor hotspots. Population density is 250 people/km^2, and the other assumptions are at base-case values (with "high" data rates).

Figure 5.16 shows the required bandwidth to achieve given throughputs for a location with 4,000 people/km^2. For this population density, 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. We also find that there 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 (Figure 5.11).

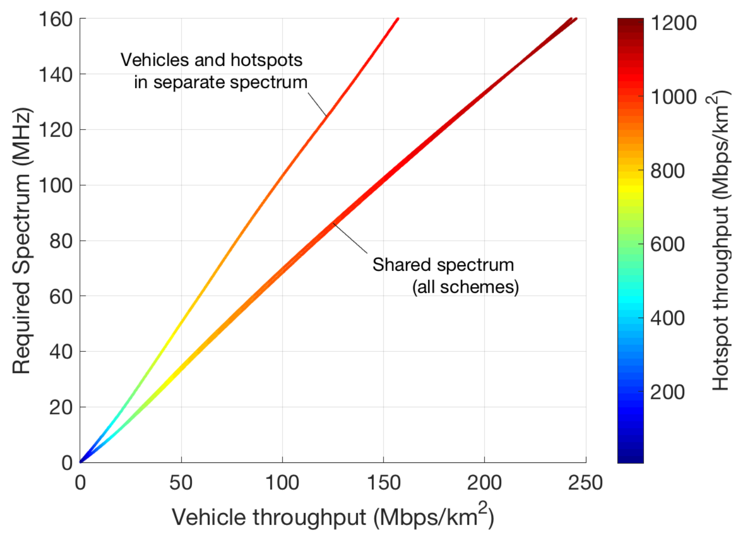

Figure 5.16. Required spectrum to achieve given vehicular and hotspot throughputs, as a function of vehicular throughput, for different strategies. Population density is 4,000 people/km^2, and the other assumptions are at base-case values (with "high" data rates).

For all scenarios examined, sharing results in less bandwidth required to achieve given throughputs when compared to V2X and unlicensed devices using separate bands. Coexistence requires the same amount of spectrum as relay cooperation. On the other hand, backhaul cooperation requires less bandwidth than coexistence. In particular, the difference is more significant with lower population densities with the presence of outdoor hotspots and is less significant in other scenarios.

5.5. Conclusions

In this chapter, we address three spectrum-related issues. The first is how much spectrum should be available for ITS. Second, whether that spectrum should be shared with unlicensed devices, as has been proposed by the FCC and others. Third, if ITS spectrum is to be shared, what sharing scheme should be implemented. For the analysis, we consider the scenario in which safety messages are transmitted over spectrum that is not shared for other types of communications. V2X and unlicensed devices may share spectrum on a co-equal basis to carry non-safety-critical information, such Internet traffic. We consider either V2X and unlicensed devices operating in separate bands, or one among three possible sharing schemes. The first is *coexistence*, where V2X and unlicensed devices sense each other transmissions, but devices of one type try to avoid interference without explicit cooperating with devices of the other type. (This is similar as one of the proposals to the FCC to share the ITS band.) In the second sharing scheme, unlicensed devices act as access points to the Internet both for unlicensed and V2X traffic, which we call *backhaul cooperation*. In the third scheme, unlicensed devices do not act as access points for V2X traffic, but rather act as part of the vehicular mesh. In this scheme, hotspots relay traffic between vehicles and RSUs or other vehicles, which we call relay *cooperation*.

On how much to allocate for ITS, we found that if spectrum is allocated exclusively, there are realistic scenarios where allocating spectrum far in excess of what is used for safety enhances social welfare, and there are also realistic scenarios where too much spectrum has already been allocated for ITS. The bandwidth that maximizes social welfare is sensitive to uncertain factors such as the penetration of devices in vehicles, data rates (particularly those to unlicensed devices), characteristics of the cellular network, and the opportunity cost of 5.9 GHz spectrum. For example, in scenarios of higher data rates and penetration, adding 40 MHz enhances social welfare if the opportunity cost is about $0.45 per MHz-pop or less. On the other hand, if data rates of Internet traffic and penetration of devices in vehicles do not reach the levels assumed, or macrocellular networks are expanded with cheaper or more efficient technologies than current cell towers, then it might be that it is not cost-effective to allocate any spectrum in excess to what is allocated for safety. Because of this uncertainty, allocating spectrum exclusively runs the risk of not providing enough spectrum for welfare-enhancing ITS.

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

One of the forms of cooperation examined in this chapter, (backhaul cooperation) has the potential to further improve the efficiency of sharing ITS spectrum, when compared to the simpler coexistence scheme. However, the magnitude of this advantage is sensitive to the conditions where sharing takes place. For example, backhaul cooperation requires less bandwidth than coexistence for lower population densities, which are representative of most of the U.S. distribution, but the difference in bandwidth is significant only if those locations have widespread presence of outdoor hotpots such those of metropolitan Wi-Fi networks. With few outdoor hotspots, the bandwidth required with backhaul cooperation is not significantly less than with coexistence. Such deployment of outdoor hotspots is unlikely for sparsely populated areas. Given that such a cooperation scheme would require regulatory efforts that are probably far more complex than a mandate for coexistence (Peha 2009), it unlikely that the benefits of cooperation outweigh the cost of implementing it. Moreover, we found that the other cooperation scheme examined in this chapter (relay cooperation) does not produce results that are significantly different from those of the simpler coexistence scheme. Therefore, a nationwide mandate of a sharing scheme such as relay cooperation over coexistence would probably not be worth the extra technical and regulatory cost.

In the recent policy debate over ITS spectrum, it has generally been assumed that the size of the ITS band would remain fixed at its current level, and the question is whether to share with unlicensed devices. If the bandwidth available to vehicles is fixed, we have found that the throughput achievable with V2X devices coexisting with unlicensed devices in shared spectrum can be significantly lower than the throughput in exclusive spectrum (up to 2/3 lower, depending on the scenario). 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 coexist, then policymakers could change regulations to increase the size of the ITS band while still giving unlicensed devices access. In other words, while unlicensed devices gain access to the ITS band, V2X devices could use the adjacent unlicensed bands for non-safety-critical traffic. (Again, sharing the ITS band might exclude the portion of the ITS band reserved for safety messages.) Under these circumstances, V2X and unlicensed devices would achieve the same throughput performance in shared spectrum while using less bandwidth overall. Such an approach would likely be implemented with the coexistence sharing scheme rather than with cooperation. While it might be reasonable to require cooperation from unlicensed devices as a condition to operate in the ITS band, the small improvement in throughput and required bandwidth (if any) of cooperation over coexistence is not likely worth the complexity of the former. Moreover, unlicensed devices do not cooperate in the bands in which they already operate, which would make cooperation in these bands even harder to implement.

Besides, we have found the throughput to unlicensed devices in shared spectrum to be not much 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.

Appendix – Comparison of required bandwidth between scenarios

We compared required bandwidth on shared spectrum with required bandwidth for vehicles and hotspots using separate bands, for scenarios other than those shown in Section 5.4. Those scenarios span a wide range of assumptions. We varied population density from 1,000 to 20,000 people per km^2, OBU penetration from 25% to 100%, data rates of incoming traffic varying from 400 kbps per vehicle and 5 Mbps per hotspot to 4 Mbps per vehicle and 27 Mbps per hotspots, and with outdoor hotspots either being present or not.

For these scenarios we find results similar with those in Section 5.4. That is, we find that less shared spectrum is required by vehicles and hotspots using shared spectrum, when compared with vehicles and hotspots using separate spectrum. In Figure 5.A we show the ratio of the total amount of spectrum used by vehicles and hotspots in separate bands, to the amount of spectrum with both types of devices using shared channels. The bars show the ratio scenarios with indoor hotspots only (*I* in the figure) compared with indoor and outdoor hotspots (*IO*), population densities (*pop*), OBU penetration (*pen*), and data rates of incoming traffic per vehicle (*Mbps/V*) and per hotspot (*Mbps/H*). For adjacent bars, one factor is varied at a time.

The graph is useful to show both the absolute magnitude of the ratios, and the difference between ratios. Figure 5.A shows that all scenarios examined have average bandwidth ratio greater than 1, meaning that shared spectrum uses less bandwidth than separate spectrum to achieve given throughputs, for numerical values that are representative of relevant ranges of assumptions. (However, for some scenarios the 95% confidence interval for the ratio suggests that the ratio can be a low as 0.75, for lower population densities and when throughput is close to the maximum achievable in the scenario.)

In particular, the graph shows that a scenario with indoor-only hotspots has a higher ratio than a similar scenario but with indoor and outdoor hotspots. This means that more spectrum is needed to achieve given throughputs when there are indoor and outdoor hotspots in shared spectrum, because of the increased interference. Comparison between indoor-only and indoor-outdoor hotspots for other scenarios of population density, penetration and data rates confirm that trend (these are not shown in Figure 5.A). Figure 5.A also shows that the ratio of bandwidth increases with data rates of incoming traffic to vehicles and hotspots. This suggests that it is worth sharing spectrum for a variation in data rates of an order of magnitude. Even if that data rates increase sharply in the future, sharing appears to be beneficial.

The bottom bars show that the ratio is similar for different population densities or OBU penetrations. The differences in ratios are not statistically significant at a 5% confidence level. However, as with most other scenarios the ratio is close to 2, which means that sharing spectrum requires as much as half the bandwidth required to achieve given throughputs in separate bands.

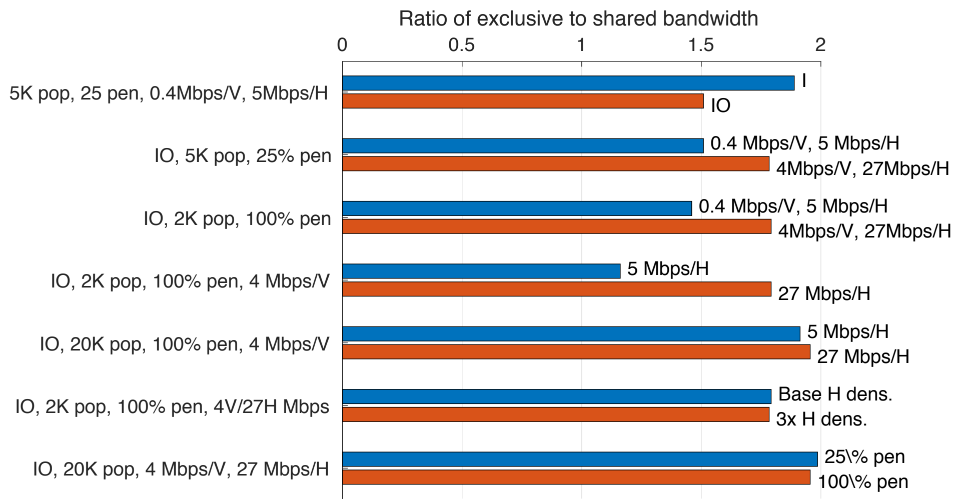

Figure 5.A. Ratio of bandwidth in exclusive channels to bandwidth in shared channels to achieve a given target throughput. The target vehicle throughput in each scenario is set as half the throughput obtained at 160 MHz. Each bar shows the ratio for a different scenario. The pairs of bars compare ratios for scenarios where one factor being is changed at a time: Indoor hotspots (I) vs Indoor and outdoor (IO), data rates (Mbps/V for vehicles, Mbps/H for hotspots), hotspot density (pop+pen), and vehicle density (pen).

Chapter 6. Conclusions

In this work we investigate the conditions in which V2X-based vehicular networks should be deployed as an important way to carry Internet traffic. Expanding the capacity of cellular networks alone may not be the most cost-effective way to meet the present and future growth of mobile Internet. We examine whether V2X-based networks should be deployed to complement cellular capacity. We address this issue by evaluating cost-effectiveness of offloading Internet traffic that would otherwise be carried over cellular infrastructure onto V2X-based networks. Depending on the scenarios being analyzed, we also evaluate other criteria as well such as cost savings for governments or spectrum efficiency. The evaluation is done under several conditions of population density, penetration of V2X devices, data rates of Internet traffic, costs of V2X and cellular devices and spectrum, bandwidth allocated for cellular networks, and spectrum efficiency of cellular technology. Moreover, the evaluation is done for varying design choices of infrastructure and spectrum. More specifically, we evaluate cost-effectiveness and other criteria with varying quantities of RSU infrastructure, with RSUs deployed either for Internet only or shared RSUs, with varying amount of ITS spectrum allocated exclusively for V2X devices, and with ITS spectrum shared with unlicensed devices. In this chapter we discuss our findings and how they apply to several scenarios.

6.1. Conclusions under a U.S. DOT mandate

In a scenario which DSRC devices are mandated in vehicles to enhance road safety, as proposed by the U.S. DOT, we conclude that it is cost-effective to deploy RSUs for Internet access. This is true initially for densely populated urban areas, but over time V2X-based networks are cost-effective in less populated areas as well. We find that benefits from Internet access alone are not enough to justify deploying devices in all vehicles, i.e. the benefits of Internet alone are less than total costs. However, this fact does not include benefits from enhanced road safety and other outcomes. These benefits are likely, and under a mandate to enhance road safety, OBUs are deployed regardless of Internet access. Moreover, the ITS band is already allocated for V2X communications. Therefore, the costs of OBUs and spectrum are incurred anyway, and V2X-based networks are more cost-effective than macrocellular networks for Internet access as long as benefit exceeds the incremental costs of deploying RSU infrastructure. We find that benefits of Internet access through V2X networks would be significantly greater than RSU cost in densely populated areas, for OBU penetrations that are representative of a few years after such a mandate becomes effective, and peak-hour Internet traffic per vehicle is compatible with current usage of mobile Internet. Some of the factors that affect benefits and costs are uncertain, such as OBU penetration, volume of Internet traffic, cellular bandwidth and spectral efficiency, and unit costs of devices. However, OBU penetration and data rates of Internet traffic are expected to increase over time. As a result, the benefit of Internet access would exceed RSU infrastructure cost in regions with lower and lower population densities over time.

Under a mandate, if RSU infrastructure is shared between local governments and ISPs for safety applications and Internet access, rather than deploying RSUs for Internet alone, then V2X-based networks would be deployed sooner in less populated areas, and in locations where it would not be cost-effective to deploy RSUs without sharing. Government agencies often deploy their own infrastructure for safety or other purposes, and they could share it with ISPs for a fee. We find that sharing is more cost-effective than ISPs deploying their own infrastructure, as long as the cost to provide Internet access with shared RSUs is less than the cost of RSUs deployed by the ISP. Moreover, governments may use the proceeds from sharing to offset their own costs as well. We estimate that governments could recover about one-fifth of the total cost to deploy safety RSUs nationwide in the U.S. Likewise, governments deploy smart city infrastructure such as streetlights with backhaul capability. They could share smart streetlights with ISPs, if they are able to upgrade them to serve as RSUs at a lower cost than RSUs for Internet only. In this case, it is even more cost-effective to share smart streetlights than safety RSUs, and nationwide savings for the government could be up to one third higher than with sharing of safety RSUs. We also examine the pricing strategies government could adopt to share. We find that the optimal prices that maximize cost-effectiveness and government savings may differ. However, maximizing government savings results in enhancements in cost-effectiveness that is also close to maximum.

In addition, if ISPs use government infrastructure at a fee that is lower than the cost of deploying their own RSUs, then ISPs would provide Internet with shared RSUs at locations with lower population densities than the locations the ISP would serve when deploying its own RSUs only. The result of infrastructure sharing is Internet at a lower cost than cellular for more of the country.

While 75 MHz of spectrum in the ITS band is allocated for V2X communications on an exclusive basis, we examine what amount is optimal from a social welfare perspective. 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

OBU penetration in vehicles, future data rates of Internet traffic, and the opportunity cost of 5.9 GHz spectrum.

That uncertainty is less relevant if ITS spectrum is shared. It is spectrum-efficient to share spectrum used by V2X-based networks with unlicensed devices, in ways similar to those proposed in response to a recent NPRM issued by the FCC. By spectrum-efficient we mean that bandwidth required to achieve given throughputs to V2X and unlicensed devices in shared spectrum is as little as half as bandwidth required when V2X and unlicensed devices use separate bands. We conclude that the spectrum available for ITS should be maintained or increased, while much of ITS spectrum should be shared with non-V2X devices.

We compare different schemes where V2X and unlicensed devices send Internet traffic on shared spectrum as equals. Those schemes differ on whether V2X and unlicensed devices coexist in a listen-before-talk etiquette, or unlicensed devices cooperate to relay V2X packets. The preferable spectrum sharing scheme should be of coexistence among equals. While spectrum efficiency could be somewhat improved by having unlicensed devices cooperate with V2X devices to carry V2X traffic, the relative benefit of such cooperation would probably not be worth the burden of mandating complex cooperation schemes over simpler forms of coexistence-based spectrum sharing.

Those findings have implications on nationwide decisions about spectrum. We show that V2X-based infrastructure for Internet access should be deployed in densely populated areas soon after a mandate for road safety is effective, and in less dense areas over time and/or if infrastructure is shared with local governments. Nevertheless, deploying V2X-based networks for Internet access is not cost-effective everywhere, e.g. in sparsely populated areas. On the other hand, spectrum decisions typically have national or regional implications. As a result, in sparsely populated areas the ITS band will be barely used. Allowing unlicensed devices use this otherwise idle spectrum at those locations would be particularly efficient. For more densely populated areas, we find that throughput to V2X devices on a given amount of shared spectrum can be significantly less than the throughput on the same amount of spectrum but used by V2X devices only. However, this finding is based on a fixed bandwidth available to those devices. If V2X devices also use adjacent unlicensed bands when spectrum is shared, then the same throughput to V2X devices achievable on spectrum used exclusively could be obtained on shared channels. There is no reason why the bandwidth used by V2X devices 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 policymakers could change regulations to increase bandwidth for V2X devices while still giving unlicensed devices access.

Our conclusions are valid for scenarios defined by several assumptions. It is important to discuss how conclusions change if some conditions differ from what we assumed. Relevant assumptions include the mandate for V2X devices in vehicles, infrastructure costs, and V2X and cellular technologies. We note these assumptions both because of their importance, and because of the uncertainty underlying them.

6.2. Conclusions with DSRC OBUs purchased without a mandate

Some of our conclusions hold for scenarios without a mandate to enhance road safety, while other conclusions are contingent on further research. All else equal, without a mandate it is not cost-effective to deploy V2X-based networks at all locations where it would be cost-effective to deploy under a mandate. However, it would probably be cost-effective in more and more locations over time. Moreover, without a safety mandate, sharing spectrum allocated to V2X with unlicensed devices is likely to be at least as efficient as with a mandate.

While in a safety mandate the only cost that matters for Internet access is the cost to deploy RSUs for that purpose, OBU costs must also be considered (at least in part) when car owners voluntarily purchase these devices for Internet access. Moreover, penetration of OBUs in vehicles is expected to grow in a much slower pace than in a mandate. If there were two locations with identical vehicle densities and V2X could be mandated in one but not in the other, after some time there should be less vehicles with OBUs in the location without the mandate, and therefore less total traffic over the V2X-based network. (However, vehicles with most usage of Internet traffic are likely the ones that have OBUs adopted first.) All else equal, the minimum population density where V2X-based networks are cost-effective should be higher when users deploy OBUs voluntarily than in a mandate. However, under a mandate we have found that over time V2X-based networks are cost-effective in less densely populated areas, meaning that these networks are worth deploying in more of the country over time. This is because data rates of Internet traffic and OBU penetration are expected to grow over time, which should be true for voluntary adoption of OBUs as well.

Moreover, sharing infrastructure deployed for safety with ISPs should be cost-effective in less locations with voluntary adoption than with a mandate, and there should be less government savings from sharing, if any. All else equal, the lower penetration of OBUs with voluntary adoption results in less governments deploying RSUs for safety applications, and those who deploy will have less RSUs. Likewise, ISPs are likely to deploy less RSUs for Internet access as well (shared or not).

However, sharing ITS spectrum with unlicensed devices is likely to be at least as spectrum-efficient with voluntary adoption as with a mandate of OBUs in vehicles. Our measure of spectrum efficiency depends on how much throughputs on shared spectrum are less than on exclusive bands. Throughput to unlicensed devices is not much less on shared than on separate spectrum. We have found this difference to be even smaller when OBU penetration is lower, which is the case under voluntary adoption. V2X throughput on shared spectrum can be significantly less than on separate bands. However, we have found this difference to be similar for a wide range of OBU penetrations as well.

6.3. Conclusions with cellular V2X technology

It is important to discuss scenarios in which V2X-based networks are deployed with new technologies either replacing or coexisting with DSRC, such as C-V2X technology. DSRC is the technology considered by the U.S. DOT in the proposed mandate, because of an almost 20-year history of development, testing and pilot deployment. However, C-V2X is a newer cellular technology in development to support vehicular communications. Although C-V2X is expected to take several years from testing to deployment, it may be an alternative to replace or coexist with DSRC in the long run, either under a mandate of devices in vehicles or not.

Some conditions are necessary for offload of mobile Internet traffic over C-V2X technology to be possible and cost-effective. One condition is that C-V2X supports throughputs compatible with Internet applications. The first specification was completed in 2017 (3GPP Rel. 14), in which the sidelink PC5 interface was intended for the exchange of safety messages only, at a rate of 10 short packets per second. The latest specification was completed in mid-2018 (3GPP Rel. 15), which allows higher throughputs to support applications such as live video over the sidelink. Offload of Internet traffic over C-V2X would be possible only with the latest release. Another condition is that the cost to deploy and operate C-V2X RSUs is less than the cost of macrocells to carry the same amount of traffic at peak hours. This depends on the coverage of those RSUs and costs of their footprint, physical structures, electronics, power and backhaul. For example, if C-V2X RSUs are mounted in small structures, with access to cheap backhaul, and with equipment that is far simpler than that of macrocells, then those RSUs may cost less than macrocells to carry given amounts of traffic at peak hours, and deploying those RSUs would be cost-effective, at least under an OBU mandate. Conversely, if RSUs have the same coverage structure, equipment and costs of a macrocell, then there is no reason to offload traffic over networks using C-V2X technology.

If the necessary conditions above are satisfied, then some of the conclusions found for V2X networks using DSRC technology may apply for C-V2X technology. The latter is claimed to have communications ranges about as twice as DSRC (Papathanassiou, Apostolos Khoryaev 2017). In this case, if the cost of a C-V2X RSU is comparable to the cost of a DSRC RSU, then the total cost of C-V2X infrastructure would be lower than DSRC RSUs. As a result, if C-V2X OBUs are mandated in vehicles, V2X-based networks would be cost-effective in locations with lower population densities than DSRC. On the other hand, with wider coverage of C-V2X governments would need to deploy less RSUs for safety applications than in a scenario with DSRC. Therefore, there would be less locations available for sharing safety infrastructure with ISPs, which would result in less benefits (cost-effectiveness and government savings) than with sharing DSRC infrastructure.

6.4. Summary

This work sheds light on the possible use of V2X-based networks to provide Internet access. These networks can offload mobile traffic, and therefore complement the capacity of cellular networks at a lower cost than expanding macrocellular infrastructure. Such offload can help serve the ever-growing demand to carry Internet traffic, especially from mobile devices. Moreover, the use of V2X-based networks for non-safety applications can help offset the investment in those networks that is necessary to enhance road safety. With this work we show that it is cost-effective to deploy V2X-based networks for Internet access under several conditions, which informs decisions about where and when to deploy V2X infrastructure and to allocate spectrum for Internet access. We also inform decisions to share infrastructure and spectrum used for V2X-based networks with other purposes, and how to implement such sharing.

Chapter 7. Future Work

There are many opportunities for further research that can be derived from our work. They include assessing scenarios that may involve carrying non-safety traffic over V2X-based networks with future technology, and scenarios with similar technology but different conditions from those assumed in our work.

One opportunity for further work is to extend our conclusions for a range of costs of RSU infrastructure wider than we assumed. The more expensive an RSU, the less cost-effective it is to deploy RSUs for Internet access when compared to expanding cellular infrastructure. While our conclusions hold even if the cost per RSU is 50% higher than the base-case value, it is worth noting that such a range a variation may be representative of RSUs deployed in places with existing mounting structures (poles or walls), with easy access to power and backhaul, and with backhaul cheaper than the cost of cellular service. This was the case of RSUs deployed in the urban area of Porto. If RSUs must be otherwise deployed in locations without such existing infrastructure, then RSU costs can be significantly higher than the range considered. Further work is needed to determine cost-effectiveness of deploying RSUs where new poles, energy and communications infrastructure must be built entirely. This can be accomplished using our engineering-economic approach described in Chapter 2.

Further work is also needed to determine the population densities where V2X-based networks are cost-effective when OBUs are deployed voluntarily rather than under a mandate, and whether this is representative of real population densities. To do this analysis with our engineering-economic approach, it is necessary to find what would be the quantity of OBUs purchased voluntarily. This quantity and the most cost-effective quantity of RSUs are mutually related. A car owner could purchase an OBU if its marginal cost is less than the marginal cost savings from offload experienced by the car owner. These cost savings depend on the quantity of RSUs deployed, which in turn depend on how many OBUs are expected to be deployed. Hence, the optimal quantities of OBUs and RSUs could be determined as solutions to a system of equations, or reaction functions, in which marginal benefit of OBU depends on RSU quantity and vice-versa. This approach can also be used to determine cost-effectiveness and governments savings of infrastructure sharing, if any, under voluntary adoption of OBUs. Upon determining the equilibrium quantity of OBUs under voluntary adoption and the respective RSU quantity, it is also possible to use the engineering-economic approach to determine spectrum-efficiency of sharing spectrum, i.e. find how much less spectrum is required by V2X and unlicensed devices on shared spectrum than on separate bands, to carry a given amount of Internet traffic.

Another opportunity for future work is to quantify cost-effectiveness, government savings from shared infrastructure, and efficiency of shared spectrum when V2X-based networks use cellular V2X (C-V2X) technology. Such work would shed light on when and where this new technology would be cost-effective to carry Internet traffic, and how it would compare with DSRC with respect to the quantified measures. One approach could be to build a network simulation model to estimate throughput capacity of networks using C-V2X and use it to replace the DSRC model in our engineering-economic method described earlier in this work.

While C-V2X can be deployed as an alternative to DSRC, a mixed scenario is also possible where both DSRC and C-V2X technologies are deployed. Automakers have already been deploying DSRC in a few car models, and transportation agencies have deployed DSRC RSUs for safety applications in a number of locations in the U.S. Moreover, the citywide DSRC network in Porto, Portugal is an example of deployment for Internet access and other non-safety applications. Hence, wide deployment of DRSC technology is a plausible scenario, whether under a mandate or not. However, after some years it is also possible that C-V2X technology has been sufficiently developed and tested such that a widespread deployment is considered by governmental or private organizations. The result could be near-term deployment of DSRC followed by mid to long-term deployment of C-V2X (Mueck and Karls 2018). It is likely that a nationwide mandate of vehicle devices would be based on a single technology, as proposed by the U.S. DOT in 2016. DSRC and C-V2X are not interoperable, i.e. devices with one technology would not be able to communicate with the other technology. Multiple technologies would also limit economies of scale that often result from a mandate. A scenario in which both technologies are used as the result of voluntary adoption is more likely, where each individual adopts the technology that is perceived to be more beneficial for him/herself with respect to safety and/or non-safety applications. (Some individuals and organizations may decide to have both a DSRC OBU and a C-V2X OBU in the same vehicle, or both a DSRC RSU and a C-V2X RSU in the same location.)

A scenario with both DSRC and C-V2X technologies is likely to result in locations where two V2X-based networks would coexist (one with DSRC and the other with C-V2X). One possibility is of two technologies without any interaction with each other, in which DSRC and C-V2X are allocated to different channels, as proposed in (5G Automotive Association 2017). In this case, two independent networks would overlap in a location. The discussion in Chapter 6 about cost-effectiveness of each technology would apply, with the difference that the penetration of OBUs in vehicles of each technology would be lower than the penetration that would be expected if a single technology is deployed. As a result, V2X-based networks would be cost-effective for Internet access in locations with population densities higher than in the scenario with a single technology. Moreover, benefits from sharing infrastructure would be lower with two independent networks than with a single technology, if any. Also, the spectrum allocated for DSRC devices could probably be shared with unlicensed devices, requiring less spectrum than allocating separate bands to DSRC and unlicensed devices. Spectrum allocated to C-V2X devices could be shared with unlicensed devices with sharing mechanisms such as Licensed-Assisted Access (LAA) for the ITS band.

It is also possible that DSRC and C-V2X coexist in shared channels, but appropriate coexistence schemes that mitigate mutual harmful interference must be developed and tested (Mueck and Karls 2018; 5G Automotive Association 2017). Ongoing work such as (Gulati et al. 2018) proposes devices with DSRC as the primary users of the band, while C-V2X devices sense the channels for DSRC transmissions. If DSRC devices are transmitting, C-V2X devices either switch to unused channels or refrain from transmitting. Further research is needed to quantify cost-effectiveness and spectrum efficiency of hybrid scenarios in which DSRC and C-V2X technologies coexist in separate or in shared channels.

While this work is about the use of V2X-based networks to provide Internet access, another direction for future research is to consider other forms of traffic. One possibility is delay-tolerant traffic such as data collected from sensors in smart-city deployments. For example, in the city of Porto sensors were deployed to monitor air quality and other environmental measurements. Data are sent to centralized servers using V2X-equipped vehicles, and most of that data is still relevant even if it takes several minutes until the sensors connect to vehicles. In this case, cellular connections can be avoided and the ratio of offloaded traffic over V2X to total sensor traffic can approach 100%. On the other hand, data rates of delay-tolerant traffic may be more or may be less than what is assumed in this work, depending on what sources of delay-tolerant traffic is considered. Therefore, an analysis of cost-effectiveness, government savings and spectrum efficiency with different assumptions for delay-tolerant traffic could inform what types of applications (other than Internet access) could help justify investments in V2X devices and spectrum.

Moreover, future work could evaluate throughput and spectrum efficiency of arrangements other than coexistence as defined in (Peha 2009), such as primary-secondary sharing. In this scheme, a certain group of devices is defined as primary and granted priority access to the channels, while other devices must defer transmissions over the shared spectrum as soon as the primary devices need to transmit[13]. Primary-secondary sharing may be considered to preserve the performance of the primary devices with respect to some quality of service requirement, while offering capacity for secondary devices on an opportunistic basis. In (Gulati et al. 2018), a sharing scheme is proposed where DSRC devices are primary users of shared channels and C-V2X devices transmit only when they don't detect any DSRC transmissions. Future work can quantity capacity and cost-effectiveness of such a primary-secondary coexistence of DSRC and C-V2X devices being used for Internet access.

New cellular technologies collectively known as 5G are expected to enhance throughput per Hz per cell, through the combined use of approaches such as massive and multi user MIMO, beamforming, more advanced modulation and carrier aggregation. In addition, 5G systems are expected to have access to more bandwidth, by using millimetre bands between 24 and 71 GHz. Both the throughput per Hz per cell and the bandwidth allocated affect the cost of carrying each unit of data over a macrocell, and therefore how much is saved when data is offloaded to V2X-based networks. All else equal, if macrocells have higher efficiency in throughput per Hz per cell, or more bandwidth available per cell than we assumed in this work, then the cost to carry data over macrocells would be lower than the cost resulting from our assumptions. On the other hand, cost-effectiveness of V2X-based networks also depend on the cost of a macrocellular tower, which may change as 5G networks employ technologies and network architecture that are different from today's cellular networks. In the long run, mobile Internet traffic is expected to grow sharply, resulting in the need to expand capacity with the mix of technologies that is most cost-effective (5G, DSRC, C-V2X, etc.). Therefore, further work is

[13] The detection of primary by secondary devices can be done by listening and decoding primary transmissions in the shared channels or some signaling medium, by centralized databases, or other mechanisms.

needed to determine cost-effectiveness of V2X-based networks to offload Internet traffic from 5G cellular networks.

Another opportunity for future work is to assess cost-effectiveness and spectrum efficiency to offload Internet traffic over V2X-based mesh networks composed not only of vehicles and RSUs, but also other entities equipped with communications devices such as pedestrians and bicycles. Deployment of DSRC technology is typically considered for vehicles and RSUs, but since it is based on the IEEE 802.11 family of specifications, there is no reason to believe it cannot be supported by mass-marketed devices such as smartphones, hotspots, sensors, etc.[14] Likewise, C-V2X is part of a set of new cellular technologies that may be incorporated into smartphones in the future. A V2X-based network that is not limited to vehicles and RSUs would carry Internet traffic from more users, such as pedestrians and cyclists. At the same time, these users' devices augment the number of relays in the mesh, i.e. they are able to relay packets for each other, and augment the number of Internet gateways, e.g. if Internet access is provided not only by RSUs but also by Internet-connected hotspots or small cells equipped with V2X technology.

Our numerical results assume that at a given location, Internet traffic is offloaded to a single V2X-based network. We also discussed the possibility of DSRC and C-V2X networks operating simultaneously, although over separate channels. Further work can assess whether two or other number of V2X-based networks from different ISPs can sustain competition in given locations. A necessary condition for competition is that operation is profitable for each competing ISP. It is possible that competition is sustainable in locations with given population density when OBU penetration reaches a certain level, but for locations with lower population densities or penetration a policy intervention for infrastructure sharing or other policy is required

[14] In Porto, Portugal, DSRC connections are used to upload data from smart city sensors for central processing (Future Cities Project 2017).

to avoid a monopoly. That analysis would require modeling the criteria of which an ISP stays or leaves a market. Such a model could be based on assumptions about revenues and costs for the ISP, and as long as revenues exceed costs it is profitable for the ISP to deploy RSUs in a location, otherwise it is not sustainable for the ISP to deploy any RSU.

BIBLIOGRAPHY

5G Automotive Association. 2017. “Coexistence of C-V2X and 802.11p at 5.9 GHz.” *5GAA Latest News*. http://5gaa.org/news/position-paper-coexistence-of-c-v2x-and-802-11p-at-5-9-ghz/.

Aittokallio, Auri. 2015. “US AWS-3 Auction Raises Record $45 Billion.” Telecoms.Com. January 30, 2015. http://telecoms.com/395312/us-aws-3-auction-raises-record-45-billion/.

Alotaibi, Mohammed, Jon M. Peha, and Marvin A. Sirbu. 2015. “Impact of Spectrum Aggregation Technology and Spectrum Allocation on Cellular Network Performance.” In *IEEE Conference on Dynamic Spectrum Access Networks (DySPAN)*. Stockholm, Sweden.

Asplund, Henrik, Andrés Alayón Glazunov, Andreas F. Molisch, Klaus Ingemann Pedersen, and Martin Steinbauer. 2006. “The COST 259 Directional Channel Model – Part II: Macrocells.” *IEEE Transactions on Wireless Communications* 5 (12): 3434–50.

AT&T. 2015. “AT&T Wi-Fi.” 2015. http://www.att.com/gen/general?pid=5949.

———. 2017. “Home Internet Data Usage FAQs - DSL Internet Support.” 2017. https://www.att.com/esupport/article.html#!/dsl-high-speed/KM1010099.

Bai, Fan, Daniel D. Stancil, and Hariharan Krishnan. 2010. “Toward Understanding Characteristics of Dedicated Short Range Communications (DSRC) from a Perspective of Vehicular Network Engineers.” In *Proceedings of the 16th MobiCom*. New York, New York, USA: ACM Press. https://doi.org/10.1145/1859995.1860033.

Balasubramanian, Aruna, Ratul Mahajan, and Arun Venkataramani. 2010. “Augmenting Mobile 3G Using WiFi.” In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services - MobiSys ’10*, 209. New York, New York, USA: ACM Press. https://doi.org/10.1145/1814433.1814456.

Balasubramanian, Aruna, Ratul Mahajan, Arun Venkataramani, Brian Neil Levine, and John Zahorjan. 2008. “Interactive Wifi Connectivity for Moving Vehicles.” In *ACM SIGCOMM Computer Communication Review*, 38:427. https://doi.org/10.1145/1402946.1403006.

Bianchi, Giuseppe. 2000. “Performance Analysis of the IEEE 802.11 Distributed Coordination Function.” *JSAC, IEEE Journal on Selected Areas in Communications* 18 (3): 535–47. http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=840210.

Boardman, Anthony E., David H. Greenberg, Aidan R. Vining, and David L. Weiner. 2001. *Cost-Benefit Analysis: Concepts and Practice*. 2nd ed. Pearson.

Boban, Mate, Tiago T. V. Vinhoza, Michel Ferreira, João Barros, and Ozan K. Tonguz. 2011.

“Impact of Vehicles as Obstacles in Vehicular Ad Hoc Networks.” *IEEE Journal on Selected Areas in Communications* 29 (January): 15–28. https://doi.org/10.1109/JSAC.2011.110103.

Calcev, George, Dmitry Chizhik, G Bo, Steven Howard, Howard Huang, Achilles Kogiantis, F Andreas, Aris L Moustakas, Doug Reed, and Hao Xu. 2007. “A Wideband Spatial Channel Model for System-Wide Simulations.” *IEEE Transactions on Vehicular Technology*.

Campolo, Claudia, and Antonella Molinaro. 2013. “Multichannel Communications in Vehicular Ad Hoc Networks: A Survey.” *IEEE Communications Magazine* 51 (5): 158–69. https://doi.org/10.1109/MCOM.2013.6515061.

Campolo, Claudia, Antonella Molinaro, Antonio Iera, and Francesco Menichella. 2017. “5G Network Slicing for Vehicle-to-Everything Services.” *IEEE Wireless Communications* 24 (6): 38–45. https://doi.org/10.1109/MWC.2017.1600408.

Cardote, Andre, Filipe Neves, Susana Sargento, and Peter Steenkiste. 2012. “A Statistical Channel Model for Realistic Simulation in VANET.” In *IEEE Vehicular Networking Conference (VNC)*. Ieee. https://doi.org/10.1109/VNC.2012.6407444.

Chang, J., G. Hatcher, D. Hicks, J. Schneeberger, B. Staples, S. Sundarajan, M. Vasudevan, P. Wang, and K. Wunderlich. 2015. “Estimated Benefits of Connected Vehicle Applications.”

Chang, Kuor-Hsin. 2015. “Wireless Communications for Vehicular Safety.” *IEEE Wireless Communications*, no. February: 6–7.

Chen, Qi, Daniel Jiang, and Luca Delgrossi. 2009. “IEEE 1609.4 DSRC Multi-Channel Operations and Its Implications on Vehicle Safety Communications.” In *IEEE VNC*. IEEE.

Cheng, Bin, Hongsheng Lu, Ali Rostami, Marco Gruteser, and John B. Kenney. 2017. “Impact of 5.9 GHz Spectrum Sharing on DSRC Performance.” In *IEEE Vehicular Networking Conference (VNC)*.

Chhaya, Hs, and S Gupta. 1997. “Performance Modeling of Asynchronous Data Transfer Methods of IEEE 802.11 MAC Protocol.” *Wireless Networks* 3: 217–34. http://link.springer.com/article/10.1023/A:1019109301754.

Cisco. 2014. “WRV210 Wireless-G VPN Router: RangeBooster.” 2014. https://www.cisco.com/c/en/us/products/collateral/routers/wrv210-wireless-g-vpn-router-rangebooster/data_sheet_c78-502735.html.

———. 2015. “Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2014 – 2019.”

———. 2016. “Cisco Visual Networking Index (VNI): North America.”

CIsco. 2017. “Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016-2021.” *VNI*. https://doi.org/1465272001663118.

Clark, David D., William Lehr, and Steven Bauer. 2011. “Interconnection in the Internet: The Policy Challenge.” In *TPRC: The 39th Research Conference on Communication, Information and Internet Policy*.

http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1992641.

Clarke, Richard N. 2014. “Expanding Mobile Wireless Capacity: The Challenges Presented by Technology and Economics.” *Telecommunications Policy*, no. 38 (December): 693. https://doi.org/10.1016/j.telpol.2013.11.006.

Comcast. 2013. “Comcast Unveils Plans for Millions of Xfinity WiFi Hotspots.” Press Release, June 10. 2013. http://corporate.comcast.com/news-information/news-feed/comcast-unveils-plans-for-millions-of-xfinity-wifi-hotspots-through-its-home-based-neighborhood-hotspot-initiative-2.

———. 2015. “XFINITY WiFi.” 2015. http://www.xfinity.com/wifi/default.htm.

Daher, Robil, and Alexey Vinel. 2013. *Roadside Networks for Vehicular Communications: Architectures, Applications, and Test Fields*. USA: Information Science Reference.

Deutsche Telekom. 2013. “Connected Cars Get Big Data Rolling.” 2013. http://www.telekom.com/media/media-kits/179806.

Engebretson, Joan. 2016. “IGR: Average Monthly Broadband Usage Is 190 Gigabytes Monthly Per Household - Telecompetitor.” 2016. http://www.telecompetitor.com/igr-average-monthly-broadband-usage-is-190-gigabytes-monthly-per-household/.

Eriksson, Jakob, Hari Balakrishnan, and Samuel Madden. 2008. “Cabernet: Vehicular Content Delivery Using WiFi.” In *MobiCom ’08 Proceedings of the 14th ACM International Conference on Mobile Computing and Networking*, 199–210. ACM New York, NY, USA. https://doi.org/10.1145/1409944.1409968.

ETSI. 2018. *Universal Mobile Telecommunications System (UMTS); LTE; Architecture Enhancements for V2X Services (3GPP TS 23.285 Version 15.1.0 Release 15)*. https://doi.org/10.1002/9780470033210.ch3.

ETSI, and 3GPP. 2018. *TS 122 186 - V15.3.0 - 5G; Service Requirements for Enhanced V2X Scenarios (3GPP TS 22.186 Version 15.3.0 Release 15)*.

European Commission. 2012. *EU Transport in Figures - Statistical Pocketbook 2012*. https://doi.org/10.2832/52252.

FHWA. 2016. “Highway Traffic Signals - FHWA MUTCD.” In *Http://Mutcd.Fhwa.Dot.Gov/Knowledge/Faqs/Faq_part4.Htm#tcsgq3*.

Flore, Dino. 2016. “The Mobile Broadband Standard.” 3GPP. 2016. http://www.3gpp.org/news-events/3gpp-news/1798-v2x_r14ietf ipwave.

Future Cities Project. 2017. “Discussing the Largest Vehicular Network in the World.” 2017. https://futurecities.up.pt/site/discussing-the-largest-vehicular-network-in-the-world-2/.

Goldstein, Phil. 2015. “In 2015, How Much LTE Spectrum Do Verizon, AT&T, T-Mobile and Sprint Have -- and Where?” *FierceWireless*, 2015.

Gulati, Kapil, Shailesh Patil, Gaurav Gupta, Sudhir K. Baghel, and Marco Papaleo. 2018.

DSRC-LTE V2V Co-Channel Long Term Coexistence. US Patent 2018/0014146 A1, issued 2018.

Hallahan, Ryan, and Jon M. Peha. 2009. “The Business Case of a Nationwide Wireless Network That Serves Both Public Safety and Commercial Subscribers.” In *37th Telecommunications Policy Research Conference*.

———. 2010. “Quantifying the Costs of a Nationwide Public Safety Wireless Network.” *Telecommunications Policy. Www.Ece.Cmu.Edu/~peha/Papers.Html*. https://doi.org/10.1016/j.telpol.2010.01.002.

———. 2011a. “Compensating Commercial Carriers for Public Safety Use: Pricing Options and the Financial Benefits and Risks.” In *39th Telecommunications Policy Research Conference. Www..Ece.Cmu.Edu/~peha/Papers.Html*. Washington, D.C.

———. 2011b. “The Business Case of a Network That Serves Both Public Safety and Commercial Subscribers.” *Telecommunications Policy. Www.Ece.Cmu.Edu/~peha/Papers.Html*. https://doi.org/10.1016/j.telpol.2010.12.006.

Harding, John, Gregory Powell, Rebecca Yoon, Joshua Fikentscher, Charlene Doyle, Dana Sade, Mike Lukuc, Jim Simons, and Jing Wang. 2014. “Vehicle-to-Vehicle Communications: Readiness of V2V Technology for Application.” U.S. DOT NHTSA.

Hazlett, Thomas W, and Michael Honig. 2016. “Valuing Spectrum Allocations.” *Mich. Telecomm. & Tech. L. Rev.* 23: 45.

Hossain, Ekram, Garland Chow, Victor C.M. Leung, Robert D. McLeod, Jelena Mišić, Vincent W.S. Wong, and Oliver Yang. 2010. “Vehicular Telematics over Heterogeneous Wireless Networks: A Survey.” *Computer Communications* 33 (7): 775–93. https://doi.org/10.1016/j.comcom.2009.12.010.

IEEE. 2010a. *1609.3 - IEEE Standard for Wireless Access in Vehicular Environments (WAVE) - Networking Services*. Vol. 2010.

———. 2010b. “802.11p-2010 - IEEE Standard -- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments.” https://doi.org/10.1109/IEEESTD.2010.5514475.

Instituto Nacional de Estatistica. 2011. “Statistics Portugal.” 2011. https://www.ine.pt/xportal/xmain?xpid=INE&xpgid=ine_indicadores&indOcorrCod=0006044&contexto=bd&selTab=tab2&xlang=en.

Katz, Raul. 2014. “Assessment of the Economic Value of Unlicensed Spectrum in the United States.” In *TPRC 42: The 42nd Research Conference on Communication, Information and Internet Policy*.

Kenney, John B. 2011. “Dedicated Short-Range Communications (DSRC) Standards in the United States.” *Proceedings of the IEEE* 99 (7). https://doi.org/10.1109/JPROC.2011.2132790.

Kerans, Andrew, Dat Vo, Phillip Conder, and Snezana Krusevac. 2011. “Pricing of Spectrum Based on Physical Criteria.” In *IEEE International Symposium on Dynamic Spectrum Access Networks, DySPAN 2011*, 223–30. IEEE. https://doi.org/10.1109/DYSPAN.2011.5936210.

Lacage, Mathieu, and Thomas R. Henderson. 2006. “Yet Another Network Simulator.” In *Proceeding from the Workshop on Ns-2: The IP Network Simulator - WNS2*. New York, New York, USA: ACM Press. https://doi.org/10.1145/1190455.1190467.

Lansford, Jim, John B. Kenney, and Peter Ecclesine. 2013. “Coexistence of Unlicensed Devices with DSRC Systems in the 5.9 GHz ITS Band.” In *2013 IEEE Vehicular Networking Conference*, 9–16. https://doi.org/10.1109/VNC.2013.6737584.

Lansford, Jim, John B. Kenney, Peter Ecclesine, Tevfik Yucek, and Paul Spaanderman. 2015. “Final Report of DSRC Coexistence Tiger Team.”

Lee, Joohyun, Yung Yi, Song Chong, and Youngmi Jin. 2014. “Economics of WiFi Offloading: Trading Delay for Cellular Capacity.” *IEEE Transactions on Wireless Communications* 13 (3).

Lee, Kyunghan, Injong Rhee, Joohyun Lee, Song Chong, and Yung Yi. 2010. “Mobile Data Offloading: How Much Can WiFi Deliver?” In *The 6th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, 1. ACM Press. https://doi.org/10.1145/1921168.1921203.

Ligo, Alexandre K., and Jon M. Peha. 2016. “Cost-Effectiveness of Using Connected Vehicle Infrastructure for Internet Access.” In *MASITE-ITSPA Annual Conference, State College, PA*.

———. 2017. “Spectrum Policies for Intelligent Transportation Systems.” In *TPRC - Research Conference on Communications, Information and Internet Policy*. https://ssrn.com/abstract=2943688.

———. 2018a. “Cost-Effectiveness of Sharing Roadside Infrastructure for Internet of Vehicles.” *IEEE Transactions on Intelligent Transportation Systems* 19 (7): 2362–72. https://doi.org/10.1109/TITS.2018.2810708.

Ligo, Alexandre K., Jon M. Peha, and João Barros. 2016. “Throughput and Cost-Effectiveness of Vehicular Mesh Networks for Internet Access.” In *IEEE Vehicular Technology Conference*. Montreal. www.ece.cmu.edu/~peha/papers.html.

———. 2017. “Is It Cost-Effective to Share Roadside Infrastructure for Non-Safety Use?” In *IEEE Vehicular Technology Conference*. Sydney: IEEE. www.ece.cmu.edu/~peha/papers.html.

Ligo, Alexandre K., Jon M. Peha, Pedro Ferreira, and João Barros. 2015. “Comparison between Benefits and Costs of Offload of Mobile Internet Traffic Via Vehicular Networks.” In *43rd Telecommunications Policy Research Conference*. www.ece.cmu.edu/~peha/papers.html.

———. 2017. “Throughput and Economics of DSRC-Based Internet of Vehicles.” *IEEE Access*

6: 7276–90. https://doi.org/10.1109/ACCESS.2017.2785499.

Ligo, Alexandre K, and Jon M. Peha. 2018b. "Spectrum for Intelligent Transportation Systems: Allocation and Sharing." In *To Appear in IEEE International Symposium on Dynamic Spectrum Access Networks, DySPAN*.

Liu, Jinshan, Gaurang Naik, and Jung Min Jerry Park. 2017. "Coexistence of DSRC and Wi-Fi: Impact on the Performance of Vehicular Safety Applications." In *IEEE International Conference on Communications*. Vol. 3. https://doi.org/10.1109/ICC.2017.7996758.

Lu, Ning, Ning Zhang, Nan Cheng, Xuemin Shen, Jon W. Mark, and Fan Bai. 2013. "Vehicles Meet Infrastructure: Toward Capacity–Cost Tradeoffs for Vehicular Access Networks." *IEEE Transactions on Intelligent Transportation Systems*. https://doi.org/10.1109/TITS.2013.2258153.

Markendahl, Jan. 2011. "Mobile Network Operators and Cooperation: A Tele-Economic Study of Infrastructure Sharing and Mobile Payment Services." KTH.

Mecklenbrauker, Christoph F., Andreas F. Molisch, Johan Karedal, Fredrik Tufvesson, Alexander Paier, Laura Bernadó, Thomas Zemen, Oliver Klemp, and Nicolai Czink. 2011. "Vehicular Channel Characterization and Its Implications for Wireless System Design and Performance." *Proceedings of the IEEE* 99 (7).

Meinilä, Juha, Pekka Kyösti, Tommi Jämsä, and Lassi Hentilä. 2009. "WINNER II Channel Models." In *Radio Technologies and Concepts for IMT-Advanced*, edited by Martin Döttling, Werner Mohr, and Afif Osseiran. Wiley.

Molina-Masegosa, Rafael, and Javier Gozalvez. 2017. "LTE-V for Sidelink 5G V2X Vehicular Communications: A New 5G Technology for Short-Range Vehicle-to-Everything Communications." *IEEE Vehicular Technology Magazine* 12 (4): 30–39. https://doi.org/10.1109/MVT.2017.2752798.

Moore, Andrew W. 2001. "K-Means and Hierarchical Clustering." In *Http://Www.Autonlab.Org/Tutorials/Kmeans.Html*. http://www-2.cs.cmu.edu/~awm/tutorials/kmeans.html.

Mueck, Markus, and Ingolf Karls. 2018. *Networking Vehicles to Everything*. De Gruyter. https://doi.org/10.1515/9781501507243-201.

Mussabbir, Qazi Bouland, and Wenbing Yao. 2007. "Optimized FMIPv6 Using IEEE 802.21 MIH Services in Vehicular Networks." *IEEE Transactions on Vehicular Technology* 56 (6): 3397–3407.

Naik, Gaurang, Jinshan Liu, and Jung Min Park. 2018. "Coexistence of Wireless Technologies in the 5 GHz Bands: A Survey of Existing Solutions and a Roadmap for Future Research." *IEEE Communications Surveys and Tutorials*, no. c: 1–22. https://doi.org/10.1109/COMST.2018.2815585.

Naik, Gaurang, Jinshan Liu, and Jung Min Jerry Park. 2017. "Coexistence of Dedicated Short Range Communications (DSRC) and Wi-Fi: Implications to Wi-Fi Performance." In

Proceedings - IEEE INFOCOM. https://doi.org/10.1109/INFOCOM.2017.8057214.

Newman, Stagg. 2008. "Testimony before the FCC En Banc Hearing, Public Safety Interoperable Communications and the 700 MHz D Block Proceeding." https://transition.fcc.gov/realaudio/presentations/2008/073008/newman.pdf.

"Ns-3 Network Simulator." 2018. Https://Www.Nsnam.Org/. 2018.

Office of Management and Budget. 1992. *Circular No. A-94 Revised*. *The White House*. https://www.whitehouse.gov/omb/circulars_a094/.

Ozer, Jan. 2016. "How Netflix Pioneered Per-Title Video Encoding Optimization - Streaming Media Magazine." 2016. http://www.streamingmedia.com/Articles/Editorial/Featured-Articles/How-Netflix-Pioneered-Per-Title-Video-Encoding-Optimization-108547.aspx.

Pang, Yimin, Joey Padden, and Rob Alderfer. 2018. "Sophisticated Spectrum Sharing Analysis: The Case of the 5.9 GHz Band." In *TPRC - Research Conference on Communications, Information and Internet Policy*.

Papathanassiou, Apostolos Khoryaev, Alexey. 2017. "Cellular V2X as the Essential Enabler of Superior Global Connected Transportation Services." IEEE 5G Tech Focus. 2017. http://5g.ieee.org/tech-focus/june-2017/cellular-v2x.

Park, Yongtae, and Hyogon Kim. 2014. "On the Coexistence of IEEE 802.11ac and WAVE in the 5.9 GHz Band." *IEEE Communications Magazine* 52. https://doi.org/10.1109/MCOM.2014.6829959.

Peha, Jon M. 2009. "Sharing Spectrum through Spectrum Policy Reform and Cognitive Radio." Proceedings of the IEEE. 2009. https://doi.org/10.1109/JPROC.2009.2013033.

———. 2013a. "A Public-Private Approach to Public Safety Communications." *Issues in Science and Technology, National Academy Press. Www.Ece.Cmu.Edu/~peha/Papers.Html*.

———. 2013b. "Cellular Competition and the Weighted Spectrum Screen." *TPRC41 Research Conference on Communication, Information and Internet Security*, 1–23. https://doi.org/http://dx.doi.org/10.2139/ssrn.2241276.

Peha, Jon M., and Fouad A Tobagi. 1996. "Cost-Based Scheduling and Dropping Algorithms to Support Integrated Services." *IEEE Transactions on Communications* 44 (2): 192–202.

Qualcomm. 2013. "Comments of Qualcomm Incorporated." *ET Docket No. 13-49*. http://apps.fcc.gov/ecfs/document/view?id=7022418821.

———. 2018. "Navigating the Rocky Roads for Safe Driving: How Can V2X Be Improved?" 2018. https://www.qualcomm.com/invention/research/projects/cellular-vehicle-communications#design.

Ramachandran, KN N., E. M. Belding, K. C. Almeroth, and Milind M Buddhikot. 2006. "Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks." In *Proceedings 25th IEEE INFOCOM International Conference on Computer Communications*. https://doi.org/10.1109/INFOCOM.2006.177.

Sandvine. 2014. “Global Internet Phenomena Report 1H.”

———. 2016. “Global Internet Phenomena - Latin America and North America.”

Santos, Adella, Nancy McGuckin, Hikari Yukiko Nakamoto, Danielle Gray, and Susan Liss. 2011. “Summary of Travel Trends: 2009 National Household Travel Survey.”

Sesia, S, I Toufik, and M Baker. 2011. *LTE - The UMTS Long Term Evolution: From Theory to Practice*. 2nd ed. Wiley.

Sheikh, Muhammad Usman. 2014. “Aspects of Capacity Enhancement Techniques in Cellular Networks.” Tampere University of Technology.

Tan, Nicholas, and Jon M. Peha. 2015. “Measures of Spectrum Holdings That Reflect Market Share and Concentration Among Cellular Carriers.” In *43rd Telecommunications Policy Research Conference (TPRC)*. Arlington, VA. www.ece.cmu.edu/~peha/papers.html.

Tian, Daxin, Jianshan Zhou, Yunpeng Wang, Yingrong Lu, Haiying Xia, and Zhenguo Yi. 2015. “A Dynamic and Self-Adaptive Network Selection Method for Multimode Communications in Heterogeneous Vehicular Telematics.” *IEEE Transactions on Intelligent Transportation Systems* 16 (6): 3033–49. https://doi.org/10.1109/TITS.2015.2422144.

Triest, Robert K. 1990. “The Relationship between the Marginal Cost of Public Funds and Marginal Excess Burden.” *The American Economic Review*.

U.S. Department of Transportation. 2015. “Planning for the Future of Transportation: Connected Vehicles and ITS.” https://doi.org/10.1016/S0001-2092(07)69737-3.

———. 2016. *Federal Motor Vehicle Safety Standards; V2V Communications - Notice of Proposed Rulemaking (NPRM)*. http://www.safercar.gov/v2v/pdf/V2V NPRM_Web_Version.pdf%0D.

U.S. Federal Communications Commission. 1999. *Code of Federal Regulations, Title 47, Part 90 § 90.371 Dedicated Short Range Communications Service*.

———. 2004. “Report And Order 03-324.”

———. 2010. “A Broadband Network Cost Model.”

———. 2013. “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 (Docket 13-49).” https://apps.fcc.gov/edocs_public/attachmatch/FCC-13-22A1.pdf.

United States Census Bureau. 2015. “QuickFacts - United States.” In *Http://Www.Census.Gov/Quickfacts/Table/PST045214/00*.

Uzcategui, R., and G. Acosta-Marum. 2009. “Wave: A Tutorial.” *IEEE Communications Magazine* 47 (5): 126–33. https://doi.org/10.1109/MCOM.2009.4939288.

Wang, Jian, Tao Wu, Yanheng Liu, Weiwen Deng, and Heekuck Oh. 2016. “Modeling and Performance Analysis of Dynamic Spectrum Sharing between DSRC and Wi-Fi Systems.”

Wireless Communications and Mobile Computing.

Wang, Zhe, and Mahbub Hassan. 2008. “How Much of DSRC Is Available for Non-Safety Use?” In *Proceedings of the Fifth ACM International Workshop on VehiculAr Inter-NETworking - VANET ’08*. New York, New York, USA: ACM Press. https://doi.org/10.1145/1410043.1410049.

Wannstrom, Jeanette, and Keith Mallinson. 2014. “HetNet/Small Cells.” 3GPP. 2014. http://www.3gpp.org/hetnet.

Wright, James, J. Kyle Garrett, Christopher J. Hill, Gregory D. Krueger, Julie H. Evans, Scott Andrews, Christopher K. Wilson, Rajat Rajbhandari, and Brian Burkhard. 2014. “National Connected Vehicle Field Infrastructure Footprint Analysis.” US DOT - FHWA. https://doi.org/DTFH61-11-D-00008.

Zeadally, Sherali, Ray Hunt, Yuh-Shyan Chen, Angela Irwin, and Aamir Hassan. 2010. “Vehicular Ad Hoc Networks (VANETS): Status, Results, and Challenges.” *Telecommunication Systems*, December. https://doi.org/10.1007/s11235-010-9400-5.

Zhang, Nandi, Marvin A. Sirbu, and Jon M. Peha. 2017. “A Comparison of Migration and Multihoming Support in IPv6 and XIA.” In *Proceedings of ISNCC. Www.Ece.Cmu.Edu/~peha/Papers.Html*.

Zhao, Xiongwen, T. Rautiainen, K. Kalliola, and P. Vainikainen. 2006. “Path-Loss Models for Urban Microcells at 5.3 GHz.” *IEEE Antennas and Wireless Propagation Letters* 5 (1): 152–54. https://doi.org/10.1109/LAWP.2006.873950.