PhD-FSTC-2015-29
The Faculty of Sciences, Technology and Communication

DISSERTATION

Defense held on 02/06/2015 in Luxembourg

to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG

EN INFORMATIQUE

by

Maximilien MOUTON
Born on 19 December 1987 in Bar-le-Duc, (France)

OPTIMIZATION OF HANDOVER PROCEDURE BETWEEN IEEE 802.11 ACCESS POINTS UNDER VEHICULAR MOBILITY IN AN URBAN ENVIRONMENT

Dissertation defense committee

Dr. Francesco Viti, Chairman
Associate Professor, Université du Luxembourg

Dr. Jérôme Härri, Vice Chairman
Professor, Eurecom, France

Dr. Thomas Engel, dissertation supervisor
Professor, Université du Luxembourg

Dr. Nicolas Montavont
Researcher, Télécom Bretagne, France

Dr. Raphael Frank
Researcher, Université du Luxembourg

Dr. Mario Gerla, Advisory Expert
Professor, University of California Los Angeles, USA

Acknowledgments

I would like to express my appreciation to my supervisors, Dr. German Castignani and Dr. Raphael Frank who has supported me over the last almost four years.

I would like to thank also my thesis director, Prof. Dr. Thomas Engel for his trust from the first day.

I wish to extend my gratitude to Dr. Nicolas Montavont, Dr. Jérôme Härri, Prof. Francesco Viti and Prof. Mario Gerla for accepting to give me the honour to be part of this Ph.D jury.

I would like to say a particularly friendly thank to my collegues PhD student, research associate and master student in particular Ismail Yildiz who helped me during six month in my researches.

Special thanks go to my wife and her familly who has always supported me.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Contents

1.1 Motivation and objectives

In the urban environment, a Mobile Station (MS) often has access to multiple networks. In addition to high-speed, high-bandwidth cellular (3G/4G) networks with wide coverage, a tissue of IEEE 802.11 networks have been deployed in the past decade, which provide fairly high bandwidth over a short range (cf. Figure 1.1). There are three kinds of IEEE 802.11 networks, with different characteristics and usages. First, *residential networks* provide a Wireless Local Area Network (WLAN) that extends personal Digital Subscriber Line (xDSL) wired Internet access. Second, *Community Networks (CNs)* are open to any member of the community, i.e., clients of the same Internet Service Provider (ISP), or users that have agreed to collaborate. Finally, *hotspot networks* are larger, outdoor deployments provided by a public authority or an ISP, which are intended to provide Internet connectivity over larger areas. Residential and community networks are characterized by their very dense, uncoordinated deployment, unlike hotspot networks where deployment is optimized to maximise coverage with a minimum number of Access Points (APs). Nowadays, almost all public urban areas are covered by IEEE 802.11 networks [Castignani 2012a, Farshad 2014].

This thesis focuses on the optimization of vehicular communications in urban environments. Although most such communications are Peer to Peer (P2P), many applications (e.g., navigation and entertainment) require that at least some vehicles have robust and stable Internet access [Gerla 2011]. In order to judiciously exploit these heterogeneous networks, an MS should take advantage of IEEE 802.11 networks and use them as a complement to, or a substitute for, other Internet access networks. As the context consists of multi-homed mobile devices that are able to dynamically select different access networks, the always-best-connected paradigm proposed by Gustafsson

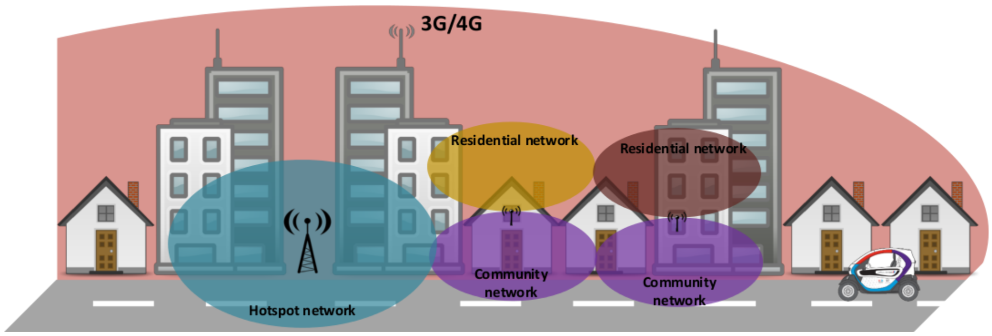

Figure 1.1: The network ecosystem of an MS in urban environment

and Jonsson [Gustafsson 2003] is relevant. In the scenario described by the authors, mobile users can select the best available access network anywhere and at any time. In order to achieve this, they describe three network mobility requirements: (i) *session continuity*, i.e., the current session between the MS and a correspondent must be ensured when moving to another network, despite the fact that the MS changes its network identity, (ii) *session transfer*, it should be possible to transfer the current application session from one device to another, and (iii) *reachability*, the MS should be reachable regardless of its current location. These requirements enable the MS to perform seamless handovers, either between Internet access networks, or between APs that belong to the same Internet access network. A handover occurs at two levels. First, at Layer-2 (L2), i.e., the MS establishes a peer-to-peer connection with the AP. Then, at Layer-3 (L3), the MS requests a network identifier (an Internet Protocol (IP) address) in order to be able communicate with other peers in the new network and elsewhere on the Internet.

In this thesis we study the L2 IEEE 802.11 handover process in order to optimize the exploitation of these networks for vehicular communications. The IEEE 802.11 handover consists of five steps: triggering, scanning, AP selection, authentication and association. In the first step, the handover is triggered based on pre-established criteria. During the scanning process, the MS sends a *Probe Request* and waits for *Probe Responses* from APs on each channel sequentially. Based on scanning results, the MS select the "best" AP according to signal strength criteria. Then the MS authenticates with the chosen AP by sending its credentials and finally, in the association process, the AP grants the MS access to the local network.

IEEE 802.11 networks are characterized by their short range. As a result, an MS must trigger frequent handovers between APs. However, this process, in its standard version, causes long disconnections. This task is even more challenging when considering vehicular communications because their mobility implies significant and frequent variation in Received Signal Strength (RSS), leading to sudden disconnections. The

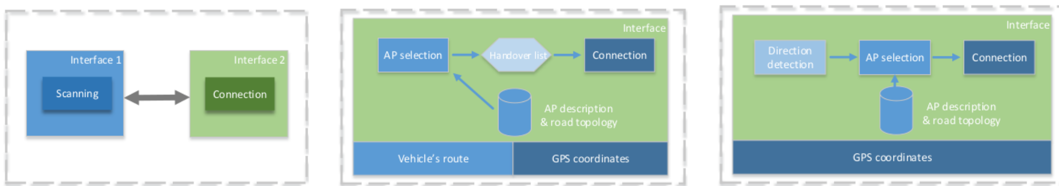

(a) Introducing a second interface (b) Bringing prior knowledge of the network characteristics (c) Including vehicular mobility prediction

Figure 1.2: Contributions description

objective of this thesis is to identify the limitations and propose optimizations for the different phases of the handover in order to eliminate L2 transition latency in the context of vehicular mobility in an urban environment. This work is a first step in the global improvement of the handover process, where the intention is to provide the MS with seamless session continuity consistent with the always-best-connected paradigm.

1.2 Thesis contribution

In this thesis, we investigate the specific characteristics of vehicular mobility and its impact on vehicular communications. First, we highlight that vehicular mobility leads to a significant instability in RSS and draw some conclusions regarding the inability of scanning to provide reliable information to the MS, leading to inappropriate AP selection. Next, we show that vehicular mobility is often predictable as it is constrained by the road infrastructure and regulations, and suggest that this characteristic could be used to improve vehicular communications.

In light of these investigations, we formulate two objectives depicted in Figure 1.2. The first is to eliminate the disconnection period created by the IEEE 802.11 handover, such that mobile users remain fully associated while in their range. The second is to optimize the AP selection process by identifying the best APs and handover location (consistent with the always-best-connected paradigm). We argue that a focus on latency reduction is not enough to provide seamless vehicular communications, because of the frequent RSS variations encountered by an MS.

This thesis addresses the two objectives presented above. Its contribution (shown in Figure 1.2) can be summarized as follows:

Introducing a second interface for seamless handovers This approach relocates the handover process, and the resulting disconnections, to a second interface (cf. Figure 1.2(a)). It is described in the literature [Brik 2005, Ramachandran 2006, Annese 2011], and addresses the first objective of this thesis. Therefore, it is presented as a preliminary approach, with the aim of evaluating: (i) the potential gain (in terms of association time) that can be achieved, and (ii) the negative impact of a scanning-based AP selection. This approach was implemented in a city-wide hotspot network. Its results were presented in [Mouton 2013a]. To our knowledge this is the first study of such an approach under these conditions.

Replacing scanning by prior knowledge of network characteristics In order to mitigate RSS instability and optimize the AP selection process, the MS needs to anticipate changes in RSS such that it is able to prevent any potential sudden disruption to the connection. The MS should also be able to detect which candidate AP provides suitable RSS (and should be selected) at an early stage. Our investigations show that the scanning process is unable to do this; thus, we consider it unsuitable for vehicular communications and argue that it can be removed.

Instead, we propose replacing it with prior knowledge of the road topology, the location of APs and a model of the RSS in nearby IEEE 802.11 networks (cf. Figure 1.2(b)). This knowledge allows the MS to forecast the RSS of the current AP in all of the directions the vehicle can take, and choose the best candidate AP and the best location to perform the handover. The results of the field tests of this approach were presented in [Mouton 2013b].

Including vehicular mobility prediction in AP selection The choice of the best AP depends on the route taken by the vehicle. Except in cases where the driver is guided by a GPS system, it is impossible to know the vehicle's route in advance. However, for optimal AP selection, it must be possible to predict the vehicle's direction. The prediction is necessary every time the driver can choose between several options.

Therefore, we propose a vehicle direction prediction system that relies on a simple analysis of the vehicle's trajectory (cf. Figure 1.2(c)). This knowledge is computed in real time and provided to the AP selection process. This, in turn assesses the modeled RSS in the predicted direction in order to select the best AP. An evaluation of this approach was published in a journal paper [Mouton 2015].

In this thesis we propose to evaluate the impact of handover triggering and AP selection based on prior knowledge of the road topology, the location of APs and a model of the RSS and compare it with legacy handover technique.

Several principles guide this research.

- **Existing standards**: The work is based on technologies that are already deployed in the field; the intention to make them usable with no delay.

- **Off-the-shelf hardware**: All network interfaces used during our experiments were commercial, off-the-shelf hardware.
- **Client-based approach**: The network protocol improvements that are proposed only imply client-side modifications, such that the solution can be applied in the context of the standard network infrastructure.
- **Field-test validation**: All the proposed solutions have been validated in the field in real-life conditions.

Despite this pragmatic approach, we were aware that future, new technologies must be taken into account in order to make the proposed solutions sustainable. We, therefore dedicated part of our work to studying the extension of our proposals in other contexts. In particular, we examined the feasibility of our solutions in IEEE 802.11p networks.

1.3 Outline

This thesis is organized as follows. In Chapter 2 we present the state-of-the-art in vehicular communications and IEEE 802.11 networks. Our proposals are outlined in detail in Chapter 3. This chapter first introduces the two objectives of the thesis, before then describing its three contributions. Chapter 4 shows the results of the field experiments performed to evaluate our contributions. Finally, Chapter 5 concludes the thesis by discussing future work and perspectives.

Chapter 2

Related work

Contents

The first part of this chapter outlines how vehicular networks are currently used, and highlights the need for communication with infrastructure networks. Next, it presents existing networks that offer access to the Internet, and some potential ways to exploit wireless network heterogeneity in the urban environment. Finally, it describes the specific case of urban IEEE 802.11 networks, and reviews the performance of vehicular communications in these networks.

The second part focuses on handover optimization (i.e. the process that allows an MS to change AP) in IEEE 802.11 networks. It describes the standard handover process and its impact on MS communications. Next, a classification of handover optimization methods is presented. This classification distinguishes between scanning and reassociation optimization approaches, and highlights the lack of information provided by scanning with regard to vehicular mobility requirements. Finally it introduces a number of handover schemes specifically designed for vehicular communications.

2.1 Connecting vehicles to IEEE 802.11 networks

2.1.1 Vehicular communications: usage and network architectures

Since the late 1990's, the development of cheap wireless network interfaces has led to the emergence of an increasing number of mobile devices. Nowadays, these wireless interfaces are an important feature of new vehicles and open up the possibility of new mobile communication paradigms. Consequently, vehicular communications has been subject of intense research over the past decade.

In their state-of-the-art review, Gerla and Klienrock [Gerla 2011] outline the main characteristics of vehicular networks and applications from a historical perspective. The authors show that new wireless technologies are unlike earlier developments in cellular phones (which became widely used in the mid 1980s), and wireless Local Access Network (LAN) networks (which began in the late 1990s). While the latter can be considered as a one-hop extension of a network infrastructure, vehicular networks find their origin in the first wireless network, Packet Radio NETwork (PRNET) that was developed by the Defense Advanced Research Projects Agency (DARPA). PRNET was designed for military purposes. There was a requirement for autonomous and portable P2P communications in order for it to be operational in any terrain. Gerla and Kleinrock state that vehicular networks are similar to PRNET in that they can be set up autonomously and spontaneously. This characteristic is particularly useful in disasters, when infrastructure is damaged or rendered inoperable. Using an autonomous (time limited) power source provided by vehicles, Vehicle Ad-hoc NETworks (VANETs) can continue to run, self-configure and transmit the critical information necessary for recovery. Beyond this extreme case, the authors highlight that P2P communications are an important dimension of vehicular communications and identify seven applications that are under active investigation:

- **Navigation safety**: collision avoidance, road conditions (construction work, accidents), etc.

- **Navigation efficiency**: traffic congestion avoidance, serious games
- **Entertainment**: multimedia, multiuser games
- **Vehicle monitoring**: Controller Area Network (CAN) bus interconnection, driver profiling, reduced carbon emissions
- **Urban sensing**: pollution, forensics
- **Social networking**: connecting drivers based on proximity or motion correlations
- **Emergency**: post-crisis communications

All of these applications rely on the concept of the Intelligent Transportation System (ITS) and Vehicle-to-Vehicle (V2V) communication is a key enabler. Vehicles would use V2V exchanges to evaluate traffic density [Nadeem 2004, Garelli 2011], avoid hazards (imminent collision with another vehicle, obstacle on the road, etc.) with very short delays [Biswas 2006, Taleb 2010], or to share files with another vehicle [Nandan 2005] that is part of the VANET. Whereas VANETs can be considered to be a member of the Mobile Ad-Hoc NETwork (MANET) family, they differ from the latter in many ways. The density and node mobility of the network in VANETs are much higher and more variable than in typical MANETs. Therefore, many elements of the network architecture, its organization and packet routing strategy have to be adapted to VANETs specificities.

Although vehicular applications rely on P2P communications, Gerla and Klienrock [Gerla 2011] argue that infrastructure also plays an important supporting role (with the exception of emergency applications). For example, it offers access to servers that store the content of entertainment and advertising applications. It also parses data collected by vehicles in order to provide the driver with useful information and enable classical Internet usage (web-browsing, email, etc.). As a result, a significant proportion of the vehicles in the network must have access to fast, short-delay and reliable infrastructure.

2.1.2 Connecting vehicles to wireless networks

The following section describes the three different technologies that can be used by a vehicle to access the Internet.

Cellular networks Cellular networks are the oldest and most widely-deployed wireless access networks. After a decade in development the first generation of such networks, Global System for Mobile communications (GSM), began to be deployed in the early 1990s. Initially they were used for cell phones. In 1995, they were extended to support data communication. The following generations of cellular networks can be identified:

- **GSM**: Data transmission in GSM networks initially used circuit-switching; this changed to packet switching based on Time-Division Multiple-Access (TDMA) spectrum-sharing via General Packet Radio Service (GPRS) and Enhanced Data rates for GSM Evolution (EDGE). GSM uses 200 kHz-wide channels in the 900 MHz, 1800 MHz, 1900 MHz bands that provide data rates of up to 384 kbps. GSM is the most widely-deployed technology, and the number of mobile connections exceeded 7 billion in 2014 [Intelligence 2014].

- **Universal Mobile Telecommunications System (UMTS)/ High Speed Packet Access (HSPA)**: Unlike GSM, both UMTS and HSPA use Code Division Multiple Access (CDMA). HSPA is an extension of UMTS and has itself been extended into Evolved HSPA (better known as HSPA+). The most recent version of Evolved HSPA uses channels between 4.2 MHz and 20 MHz wide in many bands from 700 MHz to 2.6 GHz, and provides downlink throughput of up to 168 Mbps.

- **Long Term Evolution (LTE)/ Long Term Evolution Advanced (LTE-A)**: LTE and its extension LTE-A are the latest generations of CDMA-based networks. Only LTE is currently deployed. LTE-A proposes an extension of the channel from 20 MHz to 100 MHz and downlink data rate improvements of up to 300 Mbps for LTE and 3 Gbps for LTE-A.

These cellular networks are the most widely-deployed worldwide. Their coverage in populated areas is close to 100% in western countries. They offer high throughput even at high speed. A mobile user who wants to access these networks must subscribe to an ISP. Their over-use in densely populated areas has led to restrictions being placed on the amount of data a user is allowed to download. Should this quota be exceeded, the data rate is reduced. It should be noted that data-consumption quotas have significantly increased with the emergence of LTE networks, although the latter remain sparsely deployed.

IEEE 802.11 networks IEEE 802.11 is a set a specifications for implementing WLAN. Its standardization began in the late 1990s. The initial standard has since been improved to include amendments related to the increase of throughput (IEEE 802.11a,b,g,n,ac,ad,ax), information reporting (IEEE 802.11k,u,v) or enhance transitions (IEEE 802.11r,ai,aq), for example.

The notion of a *station* defines all of the devices that can connect to an IEEE 802.11 network. The Basic Service Set (BSS) is a set of stations that can communicate with each other. IEEE 802.11 defines two modes of communication: ad-hoc or infrastructure. Ad-hoc mode provides a flat network structure based on P2P communication. In infrastructure mode, there are two types of stations: the AP and clients. The AP is the network's base station, i.e., a transceiver connecting a number of other clients to one another. Clients can be any type of mobile of fixed device. Here, the term Mobile Station (MS) is used to refer to devices embedded in a vehicle.

There are several different existing IEEE 802.11 standards released since 1999:

- **IEEE 802.11a**: Released in 1999, it uses Orthogonal Frequency-Division Multiplexing (OFDM) and several combinations of keying techniques (BPSK, QPSK, 16-QAM, 64-QAM) and Forward Error Correction (FEC) rates in 20 MHz-wide channels in the 5 GHz band that provide data rates of up to 54 Mpbs.
- **IEEE 802.11b**: Also released in 1999, it uses Direct-Sequence Spread Spectrum (DSSS) and Complementary Code Keying (CCK) in 22 MHz-wide channels in the 2.4 GHz band that provide data rates of up to 11 Mpbs.
- **IEEE 802.11g**: Released in 2003, it uses OFDM and CCK in 20 MHz-wide channels in the 2.4 GHz band that provide data rates of up to 54 Mpbs.
- **IEEE 802.11n**: Released in 2009, it introduces the Multiple-Input and Multiple-Output (MIMO) technology and Spatial Division Multiplexing (SDM) and uses up to 40 MHz-wide channels in the 2.4 GHz and 5 GHz band that provide data rates of up to 135 Mpbs.
- **IEEE 802.11ac**: Released in 2013, it uses the MIMO technology and SDM in up to 160 MHz-wide channels in the 5 GHz band that provide data rates of up to 600 Mpbs.
- **IEEE 802.11ad**: Released in 2012, it uses the single carrier mode with PSK/QAM modulation or the MIMO technology and SDM in up to 2160 MHz wide channels in the 60 GHz band that provide data rates of up to 6.77 Gpbs.

IEEE 802.11 became popular in the early 2000s. Since then, the number of domestic APs has increased and they now colonize a quarter of households worldwide (and up to 80% in some western countries) [Watkins 2012]. A significant number of these APs belong to CNs. The term CN refers to a set of APs shared by a community for Internet access. A community typically consists of people registered with the same ISP. Most of these APs broadcast far beyond the walls of individual houses into public areas and can be used by mobile devices. However, as these networks have been deployed for domestic use it is obvious that their location and transmission power are not optimized for outdoor communications. As a result, their range is limited. Moreover, multiple ISPs, sometimes supported by municipalities, have deployed large hotspot networks for nomadic users. These networks are often deployed outdoors and provide wider coverage and higher transmission power. In addition, hotspot networks can significantly reduce the latency of transitions between APs, because the MS may be not required to renew its IP configuration (IP address, default gateway, DNS) after a handover.

IEEE 802.11p IEEE 802.11p Wireless Access in Vehicular Environments (WAVE) is an amendment to IEEE 802.11 that is specifically designed to support high-speed V2V and Vehicle-to-Infrastructure (V2I) communications. Its main characteristic is that authentication and association phases are not used, as communication between mobile vehicles (or with the infrastructure) may be very short. Therefore, communication between two devices using IEEE 802.11p can start immediately. IEEE 802.11p uses channels of 10 MHz in the 5.9GHz band with doubled symbol length (8 μs vs 4 μs for other IEEE 802.11 amendments) in order to make the signal more robust to fading. In addition, the IEEE 1609.4 multi-channel architecture specifies extensions to the IEEE 802.11p MAC for multichannel operations. It defines two kinds of channels: control channels (CCH) and service channels (SCH). The control channel is use to carry control and safety messages, while all others messages are exchanged in the service channel. In order to provide Internet access to vehicles, the ITS paradigm defines the so-called Road Side Unit (RSU), which is an AP dedicated to vehicular communications. In the ITS paradigm, the MS is called an On-Board Unit (OBU)[ETSI 2004]. Although IEEE 802.11p has been designed specifically for vehicular communications and is, thus, the most appropriate access technology for V2V and V2I communications, there is currently no production deployment of this technology.

Comparison of network characteristics Table 2.1 compares the technical characteristics of cellular, IEEE 802.11 and IEEE 802.11p networks. Currently, cellular networks (UMTS, HSPA, LTE) are the most widely deployed and provide almost universal Internet access to users. In addition, since the development of new standards, they offer short latency (comparable to IEEE 802.11). However, access to these networks comes at a cost. There are still cases where mobile user fees are proportional to the amount of data consumed in a given period. In addition, even if the client's mobility is managed by the infrastructure, mobile users still experience disconnection, in particular in case of hard handover, i.e., when the old radio links in the client are removed before the new radio links are established.

IEEE 802.11 network density is high in urban environments and can provide high throughput, but their coverage is low compared to cellular networks. Although speed theoretically implies higher packet loss with IEEE 802.11 networks than with cellular networks, there is no significant impact on communications given speed limits in urban environments (up to 50 km/h)[Gass 2005].

In conclusion, the networks that can be used by a vehicle to access the Internet

	Cellular networks				IEEE 802.11						IEEE 802.11p
	GSM	UMTS / HSPA	LTE	LTE-A	IEEE 802.11a	IEEE 802.11b	IEEE 802.11g	IEEE 802.11n	IEEE 802.11ac	IEEE 802.11ad	
Standardization	1991	between 2001 and 2010	2009	In progress	1999	1999	2003	2009	2013	2012	2010
Penetration	Very high	High	Low	Not yet deployed	Low	High (urban)			Low	Not yet deployed	
Coverage	up to 35 km	1.5 km	1.5 Km	N/A	100 m to 300 m (outdoor)					~10 m	up to 900 m
Frequency	900 MHz, 1800 MHz, 1900 MHz	850 MHz to 2.1 GHz	700 MHz to 2.6 GHz	700 MHz to 2.6 GHz	5 GHz	2.4 GHz	2.4 GHz	2.4 GHz & 5 GHz	2.4 GHz & 5 GHz	60 GHz	5.9 GHz
Data rate	384 kbps	Downlink 168 Mbps Uplink 23,Mbps	Downlink 300 Mbps Uplink 75 Mbps	Downlink 3 Gbps Uplink 1.5 Gbps	54 Mbps	11 Mbps	54 Mbps	135 Mbps	600 Mbps	6.77 Gbps	27 Mbps
Bandwidth	200 kHz	4.2 to 20 MHz	1.4 to 20 MHz	1.4 to 100 MHz	20 MHz	22 MHz	22 MHz	up to 40 MHz	up to 160 MHz	2160 MHz	10 MHz
Delay	< 500 ms	from 50 to 200 ms	5 ms	N/A	0.5 to 50 ms						
Speed tolerance	Very high				Good					N/A	High

Table 2.1: Comparison of candidate networks[1]

[0] [Strobel 2013, Holma 2011, Ott 2004, Sun 2014, González 2008]
[1] [Strobel 2013, Holma 2011, Ott 2004, Sun 2014, González 2008]

are heterogeneous. In this context, the challenge is to find a connection strategy that efficiently exploits this heterogeneity.

2.1.3 Exploiting wireless network heterogeneity in urban environments

In an urban environment, a mobile device often has access to multiple infrastructure networks. In this heterogeneous context, cellular networks and 802.11 networks the most widely deployed in urban environments.

Due to discontinuity of IEEE 802.11 coverage, data offloading from cellular networks to IEEE 802.11 networks has been a natural solution to efficiently exploit this network heterogeneity[Balasubramanian 2010]. However, increasing IEEE 802.11 network density makes it possible for a vehicle to remain connected for longer periods. This may lead to a paradigm change in the definition of IEEE 802.11 networks in the urban environment. Rather than networks that can temporarily be used for delay-tolerant flows, they are now be seen as networks that can provide good Quality of Service (QoS) over the long term.

A further concern is the introduction of multihoming, i.e. the ability to connect through multiple interfaces, which has led to major issues regarding session continuity and toggling between interfaces.

Evolution of IEEE 802.11 network usage Based on the characteristics of cellular and IEEE 802.11 networks, several studies have proposed the default use of cellular networks and data offloading to IEEE 802.11 networks when possible. In [Balasubramanian 2010], the authors propose an approach that relies on the fact that although available IEEE 802.11 networks are sparse, they occur in bursts. They propose a twofold approach that schedules the transmission of delay-tolerant data in the near future and provides fast switching to cellular networks (3G) for delay-sensitive data. This solution takes as input the size of the transfer, the delay tolerance and an application-specific QoS metric. Based on these characteristics, it decides how to distribute the data across 3G and IEEE 802.11 networks. It relies on a prediction of IEEE 802.11 network throughput and a fast IEEE 802.11 to 3G transition. As APs occur in bursts, when the MS is within range of multiple APs, it is likely to quickly discover them. According to the predicted amount of data that can be offloaded through IEEE 802.11 networks, the MS decides whether to delay the sending of delay-tolerant data until it finds a potential AP or to send it directly over 3G. When the MS is connected to an AP, all data (including delay-sensitive information) transits it. Under these conditions, delay-sensitive data is more likely to experience transmission losses. In this case, the authors propose toggling to 3G when the link layer fails to deliver the packet within a delay threshold (e.g. 50 ms, proposed in [Balasubramanian 2010]).

The statistical study presented in [Watkins 2012] highlights increased AP density in urban environments. The area covered by CNs and hotspots is constantly expanding, providing Internet access to mobile users over a large part of their journey. In this context, it may be possible to use IEEE 802.11 networks as the default Internet access network. Cellular networks would only be used when there is no IEEE 802.11 network or potentially, concurrently with IEEE 802.11 networks. However, the nature of IEEE 802.11 networks (low coverage, no communication between APs) makes this change challenging. In order to provide Internet connectivity to an MS when it is in an area covered by a IEEE 802.11 network, there must be seamless transitions between APs. In other words, when the MS hands-off to an AP, the entry phase must be transparent. In particular, the most time-consuming steps must be improved: scanning and Layer-3 (L3) configuration.

Aggregating links: from multihoming to multipath The connection strategies described above imply that the MS has multiple interfaces and is multihomed. The term multihoming refers to the potential for an MS to have multiple addresses. In [Balasubramanian 2010] the MS uses one of its addresses at a time as it switches between cellular and IEEE 802.11 networks. On the Internet, a device is identified by its IP address. As a result, its identity changes over time according to the network it is connected to. However, in order to keep the current session alive, the MS needs to maintain its identity. The solution proposed in [Balasubramanian 2010] is for all applications to provide proxy support in order to keep the session alive when the MS toggles between networks. This requires establishing a proxy for each application, or in the Correspondent Node (CNo).

An alternative to this solution is to introduce a shim sub-layer between network and transport layers as in Mobile IP (MIP) or Host Identity Protocol (HIP). This sub-layer provides mobility support by distinguishing the identifier and the locator of the MS, which are currently both linked to the IP address. By storing the identity of the MS on top of its IP address, the shim sub-layer not only allows the same MS to have multiple locations, but also to transmit flows through all of them. A survey of the different protocols using the shim sub-layer is available in [Addepalli 2013].

Rather than simply toggling between cellular and IEEE 802.11 networks, a mobile user may use both networks simultaneously. The term multipath refers to this connection strategy that consists in simultaneously using multiple available networks. An example of a protocol that relies on the multipath concept is MultiPath TCP (MPTCP) [Handley 2013]. MPTCP is a set of extensions to regular TCP that enable it to run across multiple paths simultaneously. MPTCP introduces the concept of subflows (which are regular TCP sessions) managed by the sublayer MPTCP on top of them. An MPTCP session can be set up with one or several subflows. If extra paths are available, a new subflow is created and if a path is disconnected, the corresponding subflow(s) is(are) discarded. In the same manner, mobility protocols provide multipath support.

Multipath enables a broad range of connection strategies. The most obvious consists in using IEEE 802.11 networks to augment continuous connections to cellular networks. In this case, MS throughput varies from a baseline that corresponds to the throughput provided by cellular networks, to the sum of this baseline plus the throughput provided by IEEE 802.11 network, which can be very high following the latest IEEE 802.11 amendments (150Mbps with a single spatial stream). Since the characteristics of cellular networks and IEEE 802.11 networks differ, it would appear to be possible to send delay-tolerant flows through the link with the lowest latency (the remaining flows being sent through both networks at the same time). If the MS sends flows through multiple links, it can simply split a flow into two equal parts (or not) and send it through both networks. Alternatively, it can encode all or part of the flow in order to create redundancy should one of the links be unreliable, or even duplicate the transmission in both links.

2.1.4 Vehicular communications through IEEE 802.11 networks

This section has two goals: first, it outlines the main characteristics of urban IEEE 802.11 networks and their evolution over the past decade; and, second, it demonstrates network performance for vehicular communications based on tests performed between 2004 and 2014 in various North American and European cities. It highlights the challenges created when a vehicle uses IEEE 802.11 network for infrastructure access.

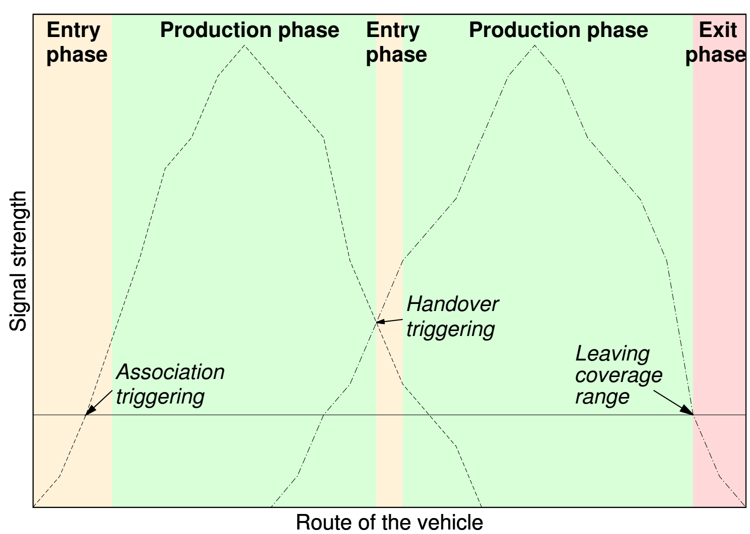

Figure 2.1: Entry, production and exit phases

2.1.4.1 Urban IEEE 802.11 network deployment

Much research has focused on evaluating the QoS available to vehicles using existing IEEE 802.11 networks, with the aim of determining the extent to which they can be used in vehicular communications. We survey the performance of urban IEEE 802.11 networks deployed over the past decade. This comparative study highlights the evolution of these networks in the past ten years and their shared characteristics.

First, we describe network deployments in terms of metrics such as AP density, channel allocation, QoS and the temporal distribution of users. Then, we focus on the impact of vehicular speed by distinguishing between the entry phase (i.e. including discovery and connection setup), and the production phase (i.e. satisfactory connectivity) of vehicular communications (see Figure 2.1) as described in [Ott 2004, Hadaller 2007]. The distinction is relevant because connectivity is often poor during the entry/ exit phases, compared to the production phase, which is mainly due to the fact that during the entry phase the MS has to initiate the connection using multiple administrative packets. It should be noted that here we consider a vehicle that opportunistically connects to all available APs (unlike the studies previously mentioned, where the scenario only concerns a single AP). The goal is to stay connected for as long as possible by initiating a connection to a new AP when it becomes a suitable alternative to the current AP. As a result, only the entry and production phases are of interest.

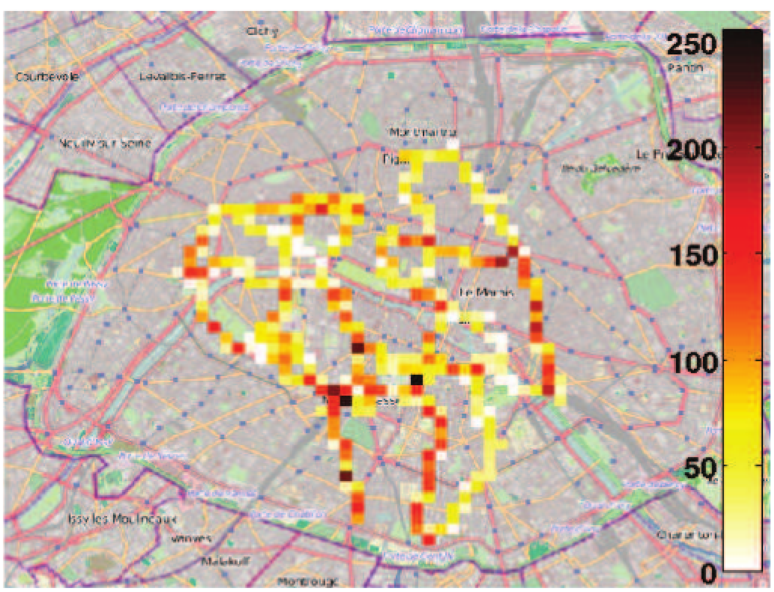

Figure 2.2: Example of the distribution of BSSIDs sensed in an urban environment [Mota 2013]

Network density Table 2.2 shows the number of BSSID (note the a given AP can broadcast one or several BSSID(s)), and coverage measured in various networks deployed between 2006 and 2014; it suggests significant growth in IEEE 802.11 network density in public areas, reaching up to 75 BSSIDs per scan. As a result, the chance of finding a suitable AP in urban areas is relatively high. Note that it would be necessary to compare the results collected in the same city at two different periods to definitively validate this trend. An example of the distribution of BSSIDs sensed is shown in Figure 2.2. In contrast, the high density of APs implies that many of them transmit their signal in the same channel, leading to potential interference (see the next paragraph for further details). In terms of coverage, the results are very different; this is due to the use of different evaluation methods. Bychkovsky et al. [Bychkovsky 2006] estimated a range of 96 m; this could be considered as the best-case scenario as the authors configured network interfaces with the maximum allowed transmission power of 200mW. The different ranges observed in [Mota 2013] and [Castignani 2011b] are explained by the definition of a lower bound of −80 dBm considered as the limit of the operative power range in the latter study. While it could be concluded that the average range of an AP is around 50 m, the acceptable range for an MS to communicate with the AP is only 30 m (i.e. 2.1 s at 50 km/h).

The development of public and private APs has meant that IEEE 802.11 networks cover a significant amount of urban public areas. Although coverage is close to complete in some locations, areas where an MS can actually access the Internet are not as extensive as for cellular networks. This is, in part, because a significant number of the available APs are encrypted. The number of open networks is in the order of 20-30% according to [Mota 2013], [Farshad 2014] and [Castignani 2011b]. Other APs belong to CNs and consequently require the user to authenticate through a Hypertext Transfer Protocol Secure (HTTPS) captive portal. However, although there are few APs that provide open Internet access, a user that belongs to a community such as Fon[2], or is subscribed to an ISP could potentially use many different APs in different cities.

Mota et al. [Mota 2013] considered three CNs that represented more than 55% of the non-encrypted APs detected during their experiments. In order to show how many APs overlap in these CNs, they built an AP graph for each CN. Two APs were considered to be connected if they were scanned at the same time at least once. The graphs of two of these CNs were sparse, meaning that they were not able to provide continuous connectivity to users by themselves. In contrast, the graph of the largest CN (27.9% of the overall number of APs) was dense, in the sense that it could potentially offer continuous connectivity to mobile users.

[2] www.fon.com

	B2006[3]	C2011[4]	M2013[5]	C2013[6]	F2014[7]
Place	Boston	Rennes	Paris	Luxembourg	Edinburgh
Mobility	Car	Pedestrian	Bus	Pedestrian	Bus
Hardware	Embedded computer	Smartphone	Smartphone	Smartphone	Smartphone
Deployment	All APs	All APs	All APs	Hotspot	All APs
# BSSID/scan min|median|max mean	0|∼2.5|20 N/A	0|22|75 N/A	1|N/A|66 27.5	0|∼2|7 N/A	0|12|45 13.4
AP coverage (m)	96 (median)	32.3 (mean)	52.35 (mean)	N/A	N/A

Table 2.2: Comparison of tested network deployments

Channel allocation and potential interference Usually IEEE 802.11 networks use 2.4 GHz and 5 GHz license-free bands. The IEEE 802.11 standard defines multiple amendments that either use one, or both (in the case of the latest amendments) (IEEE 802.11n and IEEE 802.11ac). The experiments presented in [Farshad 2014] show that currently, the 2.4 GHz band is used far more than the 5GHz band (a ratio of 10:1). This observation is surprising as the penetration of IEEE 802.11n devices is very high, but it may be explained by historical factors. IEEE 802.11g (that only operates in the 2.4G Hz band) used to be the most popular amendment, thus, to maintain backward compatibility, IEEE 802.11n APs continue to use the 2.4 GHz band.

IEEE 802.11g defined fourteen 20 MHz wide channels from 2412 MHz to 2484 MHz. Note that the 14th is only used in Japan. Of these, only channels 1, 6 and 11 do not overlap[8]. The work of [Eriksson 2008, Giordano 2010b, Castignani 2011b] shows that these non-overlapping channels represent between 78% and 83% of cases.

As mentioned before, AP density in urban environments can be very high. As most APs are distributed over three channels, this is where most interference occurs and it can significantly impact link performance. In [Castignani 2011b], a single scan showed an average of 7 APs operating in the same channel, up to a maximum of 41. In addition, APs deployed in overlapping channels, i.e. with a channel separation of less than five, interfered with each other. The number of APs affected by inter-channel interference

[3] [Bychkovsky 2006]
[4] [Castignani 2011b]
[5] [Mota 2013]
[6] [Castignani 2013]
[7] [Farshad 2014]
[8] For channels available in the United States. In Europe, other combinations are possible such as 1, 5, 9 and 13

was estimated to be approximatively 20% in [Castignani 2011b].

The impact of inter- and intra-channel interference is evaluated in [Villegas 2007], which shows that the interference generated by an IEEE 802.11 device mainly depends on the channel separation with other devices (AP or MS). The worst case is observed with a separation of one, two, zero and three channel(s) respectively (with between 25% and 75% of packets lost). With four channel separation, the impact of interference on received throughput is not significant. Finally, five channel separation has almost no impact on link quality. In [Niculescu 2007] the authors define three ranges: the communication range where two nodes can decode a given fraction of frames; the carrier sensing range where two nodes can sense each other but cannot send frames; and the interference range where two nodes cannot sense each other but can still interfere. The interference range leads to the notion of a hidden terminal or interferer which negatively impacts its neighbours' communication, while it remains undetectable. Note that the interference described in [Villegas 2007, Niculescu 2007] was evaluated in an indoor environment (offices or labs).

One of the side effects of inter-channel interference is that an MS is able to scan APs that are operating outside the current channel. The probability that such an event occurs is estimated to be 0.38 in [Castignani 2011a].

RSS distribution The RSS distribution in an urban deployment reflects the "true" coverage of IEEE 802.11 networks, i.e., the area where an MS can communicate with the AP without high packet loss.

Experiments performed in urban environments in [Castignani 2012a, Farshad 2014] have shown that RSS is usually distributed over a range of 25 dB. The distribution differs according to the device used to sense the network. When using a smartphone (e.g. the Samsung Galaxy S, Galaxy S3 and HTC Nexus One in the experiment cited here), the lower 15 dB limit represents between 80% and 90% of the sensed RSS and the median is under −80 dBm. In contrast, the RSS distribution is smoother and the median value is around −75 dBm when the same experiments were performed using a laptop, probably because the laptop was equipped with a high gain antenna.

The correlation between RSS and throughput has been the subject of much research. Zhong et al. [Zhang 2008] state that hardware-based RSS measures are not often calibrated, and that packet delivery predictions based on RSS are often overestimates. However, although the RSS distribution is sometimes not a good indicator of the QoS provided by an IEEE 802.11 network, [Aguayo 2004, Mhatre 2006, Castignani 2012a] confirm that there is a correlation between RSS distribution, Signal-to-Noise Ratio (SNR) and throughput received by the MS (similar to the packet loss rate). These studies also showed that there is a RSS threshold below which throughput begins to degrade and packet loss increases. These observations allow us to define empirical criteria that an IEEE 802.11 network must meet if it is to provide a suitable QoS to an

MS.

2.1.4.2 Vehicular communications using IEEE 802.11 networks

This section focuses on the impact of vehicular mobility on MS communications using IEEE 802.11 networks. First, we study the impact of speed on IEEE 802.11 communications. Then we review various studies that evaluate the production and entry phases using IEEE 802.11 networks. As described earlier, the production phase as defined in [Ott 2004, Hadaller 2007] and corresponds to the period when the MS has reliable Internet access. It ends when the link quality deteriorates or when the MS triggers a new AP discovery phase and toggles to a new entry phase. The term ”entry phase“ refers to the entry into the range of an AP. Both studies observed the connection between one MS and a single AP. Here, we consider the scenario where an MS opportunistically connects to dense IEEE 802.11 networks. In this case, the entry phase defines the handover, i.e., the time that elapses between the start of AP discovery and the first packet that is successfully exchanged between the MS and the Internet.

Effect of speed on link quality Gass et al. [Gass 2005] performed a set of experiments to quantify link quality using an IEEE 802.11b AP at different speeds (ranging from 5 to 75 mph or 8 to 120 km/h). The study looked at the network range, associated time, association time and packet loss, but did not find any significant impact of speed on the signal strength received by an MS even at relatively high speeds. Bychkovsky et al. [Bychkovsky 2006] found similar results. Due to the limited range of IEEE 802.11 APs, the association time is short. In both studies, the authors pointed out that this phenomenon is critical for vehicular mobility. When the vehicle speed reaches 50 km/h, an MS remains associated with an AP for only a few seconds. In these conditions, a smooth handover becomes critical in ensuring the association time is a long as possible.

Evaluation of connection time Table 2.3 is a comparison of connection time, i.e., the duration of the connection to the Internet, based on six studies published between 2006 and 2013.

Table 2.3 shows that these times are of the same order (between 10 s and 27.5 s) for all studies except one. This could be explained by two factors. First, the range of APs in the different types of networks (IEEE 802.11b/g/n) is roughly the same. As mentioned earlier, most networks operate at the same frequency (2.4 GHz). In addition, the same class of hardware with roughly the same sensitivity was used in all experiments (except for the last two). The network configuration allowed layer-2 roaming (the MS keeps the same IP after a transition) only in [Deshpande 2010] and [Castignani 2013]. Otherwise, the connection was limited to the coverage of one AP. Note that in [Deshpande 2010] there is no significant gain in terms of connection time, contrary to [Castignani 2013].

However, the disconnection time shown in Table 2.3 suggests that the network that was tested in the latter study was much denser. Table 2.3 also shows a significant increase in throughput. This could be explained by the emergence of new IEEE 802.11 amendments and high AP density. Outdoor hotspots, included in [Castignani 2013] do not correspond to high throughput due to administrative restrictions. In contrast, disconnection time was significantly lower in the more recent studies, despite an increase in APs requiring authentication through captive portals, which significantly delay the connection. This change in disconnection time highlights that the density of IEEE 802.11 networks has significantly increased (which confirms the findings reported in Section 2.1.4.1) allowing the MS to find a suitable AP more easily.

These studies show that a moving vehicle that opportunistically connects to available IEEE 802.11 networks can enjoy high throughput for a significantly long period. In addition, most disconnections correspond to a handover. Therefore, it seems that it will soon be possible to use IEEE 802.11 networks to continuously connect an MS in urban areas assuming that the remaining white areas (i.e. areas where there is no IEEE IEEE 802.11 connectivity) can be covered and if disconnections due to AP discovery and transitions can be overcome.

	B2006[9]	E2008[10]	B2010[11]	D2010 Short[12]	D2010 Long[13]	C2012[14]	C2013[15]
Place	Boston	Boston	Amherst	New York	New York	Rennes	Luxembourg
Mobility	Car	Taxi	Bus	Car	Car	Pedestrian	Pedestrian
Hardware	Embedded computer	Embedded computer	Embedded computer	Laptop	Laptop	Smartphone	Smartphone
Deployment	All open APs	All open APs	All open APs & EMN[16]	Indoor & outdoor hotspot	Indoor & outdoor hotspot	All open APs	Outdoor hotspot
Connected time (s)	24 (mean)	10 (mean)	∼10 (median)[17]	∼26 (mean)	∼11 (mean)	27.5 (median)	132.5 (mean)[18]
Disconnected time (s)	260 (mean)	126 (mean)	N/A	∼10 (mean)	∼15 (mean)	5 (median)	N/A
Downstream throughput (kBps)	30 (median)	95 (mean)	35 (mean)	300 (median) mode around 340	62 (median) mode around 340	∼60 (mean)	110 (mean) [19]

Table 2.3: Comparison of tested network deployments

[9] [Bychkovsky 2006]
[10] [Eriksson 2008]
[11] [Balasubramanian 2010]
[12] [Deshpande 2010]
[13] [Deshpande 2010]
[14] [Castignani 2012a]
[15] [Castignani 2013]
[16] experimental mesh network
[17] complementary results presented in [Soroush 2012]
[18] the session survived multiple handovers
[19] the throughput was administratively shaped to 2mbps (250kBps)

	[Bychkovsky 2006]	[Eriksson 2008]	[Castignani 2012a]
Place	Boston	Boston	Rennes
Mobility	Car	Taxi	Pedestrian
Hardware	Embedded computer	Embedded computer	Smartphone
Deployment	All APs	All APs	All APs
Scanning	750 ms	306 ms (expected)	2 s
Association	560 ms	25 ms	median
Dynamic Host Configuration Protocol (DHCP)	1800 ms	197 ms	3 s
L3 connection delay	N/A[20]	∼35 ms	median

Table 2.4: Comparison of handover phases

The impact of handovers Table 2.4 shows the duration of the various steps in the handover, which is composed of:

1. **Scanning**: the MS probes nearby APs and selects the best one.
2. **Association**: the MS sends authentication and association requests to the AP in accordance with IEEE 802.11 standards.
3. **L3 configuration**: the network grants an IP configuration to the MS through DHCP.
4. **L3 connection delay**: the time that elapses before receipt of the first packet from the Internet. It includes the Address Resolution Protocol (ARP) and Domain Name System (DNS) request/ response exchange and the sending of a packet to a test server.
5. **Authentication to a CN**: through a captive portal (not always required).

Scanning and association correspond to the layer-2 handover, while the other steps are related to the layer-3 handover and above.

The results of the three studies presented in Table 2.4 show that the time required for the handover is in the order of several seconds. The difference between the ∼3 s in [Bychkovsky 2006] and the ∼5 s in [Castignani 2012a] could be explained by the extra time spent by the MS for captive portal authentication. On the other hand, [Eriksson 2008] attempted to optimize the first three steps of the handover. Scanning

[20] The results suggest that the layer-3 connection delay lasts a few seconds.

optimization consists of scheduling each channel's probes according to the probability that there is an AP in this channel (more details in Section 2.2.2.1). Association is optimized by sending the association request immediately after the authentication request (without waiting for the authentication response). Finally, the IP configuration retrieval using DHCP is optimized by reducing timeouts. This results in a significant reduction in the overall duration of the handover although it is still too long to provide a seamless connection between two APs.

In all of these studies, retrieving the IP configuration through DHCP is one of the most time-consuming steps, and there are few workable solutions. In addition to the method proposed in [Eriksson 2008], an IP configuration caching system is described in [Bychkovsky 2006]; however, this solution has a limited impact on DHCP delay. The emergence of IPv6 may reduce the time for IP configuration thanks to the Stateless Address Autoconfiguration (SLAAC) [Narten 2007].

The layer-2 handover is also slow, particularly the scanning phase, which needs to be optimized. As shown in the previous section, IEEE 802.11 networks can potentially provide a good QoS to vehicles over long periods. The challenge is to speed up the layer-2 handover in order to make it transparent for the mobile user.

2.2 L2 handover optimization

2.2.1 IEEE 802.11 discovery and association

Description In IEEE 802.11 networks the handover takes place client side and is composed of five phases (see Figure 2.3):

- **Triggering**: The MS decides when to start the handover process.
- **Scanning**: The MS probes for APs in the different channels.
- **AP selection**: The MS select the "best" AP according to signal strength criteria.
- **Authentication**: The AP checks the MS's identity.
- **Association**: The AP grants the MS access.

The IEEE 802.11 standard defines scanning and authentication/ association procedures.

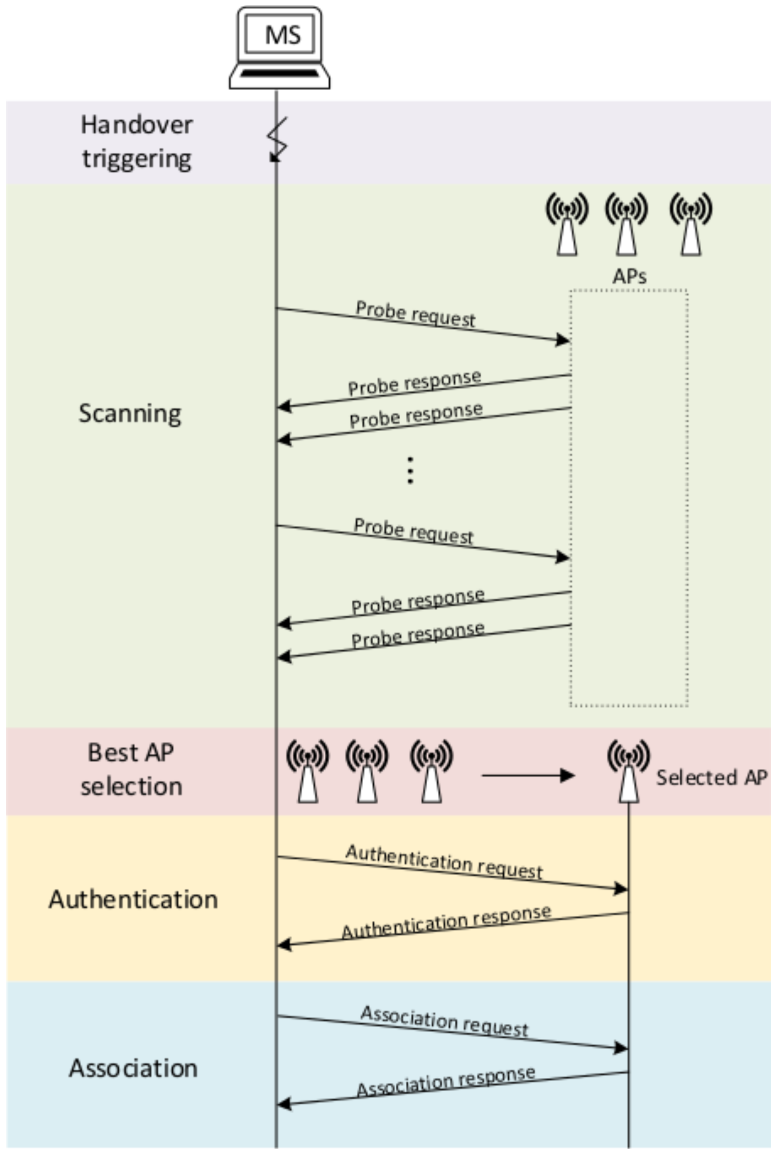

Figure 2.3: The handover process in an open IEEE 802.11 network

Scanning can be either passive or active. In passive scanning, the MS checks each channel in turn and waits for periodic beacons that are broadcast by nearby APs. In active scanning, the MS sends a *Probe Request* and waits for *Probe Responses* from APs on each channel sequentially. The time spent probing a channel is defined in IEEE 802.11 standard by two timers. Minimum Channel Time (MinCT) is the minimum time the MS must wait for a response to a probe on a channel. If at least one AP replies within MinCT, the MS will remain on the channel until Maximum

Channel Time (MaxCT) has expired. Once the MS has selected an AP from the APs that were discovered during scanning, it starts the association procedure. This begins with the authentication phase in which the identity of the MS is verified by any of the protocols used by the AP. If the AP is open, authentication is limited to an *authentication request/response* message exchange. It ends when the AP grants network access to the MS. This is done by the exchange of *association request/response* messages.

Although there is no specific rule to trigger the handover, network managers are usually configured to stay connected to the same AP until the link is disconnected, as they are not designed to handle mobility. Similarly, there is no specific rule with respect to AP selection and the usual solution is to select the AP with the highest RSS.

In the remainder of this manuscript, the term *scanning* will refer to *active scanning*, as described above. As the range of IEEE 802.11 APs is short (see the description in Section 2.1.4.1), MS handovers between APs are frequent. During the handover, the MS cannot send or receive data frames, because the wireless interface does not work on the same channel as the current AP. In practice, this means that the MS is disconnected during the whole handover process. The duration of the handover depends on multiple factors, such as the number of channels to scan, channel timers, the time taken for AP selection, or the packet loss rate (high packet loss increases the association time due to retransmissions).

The previous section showed that the native handover scheme is inefficient for outdoor mobility. In all of the studies presented, the handover process takes several seconds. Prior studies of the standard IEEE 802.11 handover mechanism [Park 2004, Mishra 2004, Velayos 2004, Teng 2009, Castignani 2011a] identified the scanning phase as the most costly, accounting for up to 90% of the time taken for the whole handover procedure.

The following sections present several handover optimizations that have been proposed over the past decade, in particular regarding scanning. These enhancements have been designed for different mobility contexts (indoor, pedestrian, vehicular, etc.).

2.2.2 Scanning optimization

As mentioned before, scanning is the most time-consuming phase of the handover. Consequently scanning optimization has been the focus of a large number of studies into handover optimization. Here, we present a survey of the most relevant scanning optimizations.

2.2.2.1 Selective scanning

Kim et al. [Park 2004] propose an extension to the Neighbor Graphs (NGs) used to optimize the association described by Mishra et al. [Mishra 2004], by including data about the AP channel. This extended NG structure optimizes scanning in two ways. First, the MS only needs to scan channels where there is at least one AP and second, because the MS knows which channels contain APs, it can switch to the next channel when it has received all of the *probe responses* it expects rather than waiting for MaxCT to expire. In order to provide NG data to the MS, the authors propose storing them on an NG server. This solution has been tested in a testbed composed of two to four APs. Under these conditions, the duration of the handover is reduced from 320 ms to an average of 12 ms to 30 ms.

In practice, many studies show that the vast majority of APs operate in three non-overlapping channels: 1, 6 and 11 [Eriksson 2008, Giordano 2010b, Castignani 2011b, Farshad 2014]. As a result, scanning can be speeded up by only probing the channels that are likely to be used by a nearby AP. In [Shin 2004] the authors introduce the following selective scanning algorithm: Initially, the MS performs a full scan of all channels. A channel mask is set by turning on the bits for all the channels in which an AP responds. Next, the best AP is selected from among those that were probed. The channel of this AP is removed from the channel mask and the three non-overlapping channels are added, given the higher probability that a probe will succeed. If there is no response to scanning the channel list, scanning begins again in the channels that were not previously selected. If there is still no result, a full scan is triggered. This approach is efficient when AP density is low. However recent studies show that in urban environments it is often high (see Section 2.1.4.1). Even if most APs use one of channels 1, 6 or 11 the chances of finding an AP in other channels remains high. Consequently, an MS using the algorithm tends to have longer scanning times as AP density increases.

In [Eriksson 2008] the authors proposed a scanning schedule whereby a given channel is allocated a scanning rate that depends on the probability that there is an AP in it. Scanning ends as soon as a suitable AP is found. The authors state that on average the expected number of scans before finding an AP is 3.64, however this number depends on AP density. As scanning ends when the first AP is found, the approach gives only a partial view of the MS's environment which may lead to the selection of an inappropriate AP.

2.2.2.2 Periodic scanning

Another approach consists of spreading channel scanning over time by dividing it into multiple subsets in order to split and dilute disconnection time over a longer period. This approach has been proposed by Montavont et al. [Montavont 2005] and Liao et al [Liao 2006]. In [Montavont 2005], the MS periodically scans a channel and lists discovered APs. The periodicity of the scanning changes depending on the urgency of finding an AP. If the RSS of the current AP is high (higher than −75 dBm) or a candidate AP has already been found, the interval between scans is set to a random value between 1 s and 2 s. In contrast, if the RSS of the current AP is low, scans are triggered at random intervals ranging from 200 ms to 300 ms. The selection is performed as follows: the MS takes the last entry on the candidate AP list and attempts to connect to it. After three failed attempts, it is considered out of range and the MS attempts to connect with the next AP in the list. In [Liao 2006], the MS returns to the current channel for a given length of time between each group scan in order to send data frames that could not be transmitted during scanning. In order to shorten the scanning phase, the authors propose a reassociation attempt if the MS finds a suitable AP after scanning a group of channels without scanning the rest. As the scanning process lengthens, the handover is triggered earlier by defining a high RSS threshold. When the RSS drops below this threshold, the MS triggers scanning. If all available APs' RSSs are lower than this threshold, the threshold is decreased and another scan is triggered. If the MS still cannot find any APs with RSS above a given minimum threshold, it selects the best candidate AP. This approach reduces packet delay, jitter and packet loss. In other words, it makes the handover smoother. Packet loss is reduced because the MS stops scanning when a good candidate is found, and buffer overflow is limited as the MS transmits data frames buffered during the scan between each channel sub-scan. The time interval between each group scan is estimated according to the mobility of the MS. For instance for a MaxCT of 50 ms, an interval of 50 ms leads to a maximum scanning duration of 1.1 s which is appropriate to an indoor mobility scenario (i.e. walking speed).

2.2.2.3 Synchronizing beacons

The IEEE 802.11 standard requires APs to periodically broadcast a beacon frame. Typically, this happens at fixed intervals of 100 ms. Ramani et al. [Ramani 2005], propose SyncScan, a modification to the infrastructure where all APs in the same channel broadcast beacon frames at the same time. Moreover, the schedule of group broadcasts (corresponding to all the APs in a channel) is staggered across channels. Clients can passively scan by switching channels when a beacon is about to arrive. This solution has several advantages. First, scanning can be spread across beacon listening periods. This significantly reduces latency as, when a handover must be triggered, the MS can immediately choose a candidate AP and switch to the authentication/ association phases. In addition, periodic scanning improves AP selection as the MS discovers candidate APs with good RSS before the RSS of the current AP becomes too low. However, this solution requires that APs have accurate clocks and use a time synchronization protocol. Furthermore, the authors highlight another issue caused by synchronization. If there are multiple APs in the same channel and they broadcast their beacon at exactly the same time, there will be signal collision, although the problem could be mitigated by simply adding a random transmission delay to each AP. Finally, frequent channel switching, even for a short period, implies a per-scan overhead expressed in Equation 2.1.

$$Overhead = ChannelNum * (2 * SwitchTime + WaitTime) \tag{2.1}$$

Where $ChannelNum$ is the number of channels to be scanned, $SwitchTime$ is the time the network interface takes to switch channel, and $WaitTime$ the time the MS waits for beacons. Given a potentially large number of available channels, this overhead could be significant.

2.2.2.4 Reducing scanning timers

The IEEE 802.11 standard does not suggest any value for MinCT and MaxCT. Consequently, IEEE 802.11 interface drivers are free to set any value, resulting in different scanning durations. If the setting is too conservative, the MS may have to wait for a long time in order to ensure that all nearby APs are found. Therefore, reducing the timer is one option that may shorten the handover. On the other hand, the modification may reduce AP detection efficiency, as a low value of MinCT may be too short for the first AP to respond and lead to the mistaken conclusion that the channel is empty. In the same way, a low MaxCT may not allow all APs to respond in time. In this case, the MS only obtains a partial view of the network environment.

In [Velayos 2004] the authors proposed theoretical best values for MinCT taking into account optimal channel conditions (the channel is idle and propagation time and probe response generation are not considered). According to the Distributed Coordination Function (DCF), under these conditions, MinCT is the sum of the DCF Interframe Space (DIFS) and the maximum number of slots in the minimum contention window, multiplied by the length of a slot. This results in a minimum value of $670\mu s$. The optimal value of MaxCT is obtained by simulation. Although their results show that this value is not bounded because it depends on the number of stations operating in the same channel, the authors suggest 10 TU (1 TU = $1024\,\mu s$) in order to avoid overlapping responses from APs.

The empirical values for AP response time reported in [Castignani 2011a] are fairly different from the theoretical results presented earlier. The authors' observations led them to conclude that AP response time varies with the hardware used and that discovery time increases proportionally with the number of APs operating in the same channel. They found a minimum mean time of 1.28 ms for indoor deployment. However, outdoor deployment showed significantly higher response times (mean values up to 8 ms). Again, channel concurrency is the main explanation for this observation. In addition, the authors show that the AP response inter-arrival time is stable for up to six responses at around 3 ms on average. This means that timers can be set dynamically according to the overall target number of responses. The results of this study tend to suggest a minimum value for these timers of around 10 ms due to the risks mentioned earlier. This study showed that the overall scanning duration can be significantly reduced, compared to the results found in urban networks discussed in Section 2.4.

2.2.2.5 Introducing a second network interface

Another interesting way to achieve fast and smooth handovers in IEEE 802.11 networks uses a second radio. Brik et al. [Brik 2005] and Ramachandran et al. [Ramachandran 2006] describe two strategies that can by applied using this principle. A first, naive approach consists in giving a fixed role to each interface. In other words, one interface is assigned the scanning role and the other is assigned the data-connection role. The scanning interface continuously scans nearby APs and, at the end of each scan, feeds the information to the MS that performs AP selection and triggers the handover. The MS then informs the data-connection interface that it must handover to the candidate AP. This approach significantly improves handover latency as the data-connection interface only performs the authentication/ association phase.

A second, more efficient approach consists in allowing the scanning interface to associate with the candidate AP while the data-connection interface is still connected to the current AP. If there is a successful connection, the interfaces exchange roles. In this approach, the whole handover process (and its cost) is transferred to the second interface and the data-connection interface stays reachable during the transition. In order to keep the same layer-3 configuration for MS handovers within the same IEEE 802.11 network, the role transition between interfaces should be transparent for the corresponding APs. To achieve this, both interfaces must have the same Medium Access Control (MAC) address. Experimental work consists of a comparison of the cross correlations between the original and the received audio signal in a Skype[21] call in [Brik 2005] and an evaluation of the inter-arrival time and packet loss in [Ramachandran 2006]. The results of these experiments reveal that the handover is completely seamless with the second method. Although these studies provide a real-life proof of concept, their relatively simple testbeds that rely on short-distance indoor mobility do not fit vehicular mobility patterns. More recently, Annese et al. [Annese 2011] proposed BATMAN, which is a routing protocol designed for vehicular communications. The layer-2 implementation of this protocol provides seamless connectivity during a handover using a second radio but focuses mainly on layer-3 issues. In the same context, Ivov et al. [Ivov 2005] propose an optimization of MIPv6, opening up the possibility of using two interfaces simultaneously. The MS uses the second approach described above. Here, when the second interface is connected to the candidate AP, and in order to increase the chance

[21] http://skype.com

that the MS receives all the packets from the HA, the authors propose that the HA bi-casts the stream to both interfaces. Bi-casting lasts until the RSS received from the new AP is high enough. Their simulation of the duration of the reconnection process shows that the approach is suitable if both the current and the candidate AP are reachable within a period of at least 510 ms.

2.2.2.6 Using RSS trends

In order to provide dynamic knowledge of nearby APs, Mhatre et al. [Mhatre 2006] propose exploiting the broadcast nature of the wireless medium by performing continuous passive scans in the current and overlapping channels (i.e., taking into account all incoming beacons). Based on this continuous information, the MS computes smoothed RSS trends using the Exponential Weighted Moving Average (EWMA) filter. Then, it predicts the RSS in the next time interval based on these smoothed trends using the Least Squares Estimator (LSE). This means that the AP selection process does not select the best candidate at a given time; instead it selects the AP that will durably outperform the current AP (based on short-term RSS changes). As a result, the handover paradigm changes from a negative statement: avoiding poor connectivity and packet loss; to a positive one: selecting an AP that is more likely to provide good RSS in the near future. A similar approach is proposed by Sadiq et al. [Sadiq 2012].

Nevertheless, the capability of such an approach to select the most suitable AP is limited as the probability that a scan finds an AP operating outside the current channel is only 38% in [Castignani 2011a] and the maximum channel distance observed is three. As a result, the MS only has a partial view of the environment, which significantly reduces the chances of finding the most sustainable AP. In addition, this kind of prediction is only useful in the short term as it does not take into account any information about the user's motion. This is a critical issue when the MS is in a vehicle. For instance, the MS will not be able to predict a sudden signal decrease when the vehicle leaves the Line of Sight (LoS) of the AP after turning at an intersection.

2.2.2.7 Scanning avoidance

In addition to the interleaved scanning strategy proposed by [Montavont 2005], the authors suggest that the MS builds a list of the characteristics of target APs (MAC address, channel and SSID) and attempts a positive association (or reassociation) request. In this scenario, the MS can skip the discovery phase described earlier, and instead uses the list of AP targets. Another approach that attempts to avoid scanning is proposed in [Khan 2012b] using the IEEE 802.21 Media-Independent Handover (MIH) standard. The standard defines data structure and communication resources intended to help MIH users in network discovery.

IEEE 802.21 is usually used to manage vertical handovers from one type of access network to another, but it can also be used for transitions between the same type of network. The handover procedure is started when the MS receives the MIH Link Going Down (LGD) signal from Layer-2 (L2). At that stage the MS sends a MIH Get Information Request (GIR) message to the Media Independent Information Service (MIIS) including its current GPS location and a custom field that tells the MIIS that the MS is looking for the identity of the next AP (Basic Service Set ID (BSSID) and channel). Upon reception of a GIR, the MIIS computes, for all nearby APs, a fitness score based on the distance from the MS, the current load and the authentication method. Next, the MIIS communicates the selected AP to the MS, which can either directly attempt a connection, or scan its channel in order to make sure it is operating (as described in [Khan 2012a]).

Another infrastructure-based AP selection scheme is proposed in [Berezin 2011]. This approach relies on the concept of the Virtual Access Point (VAP) presented in [Grunenberger 2010]. The AP selection process starts when the AP detects that the RSS of the MS has dropped below a given threshold. The current AP starts looking for an alternative AP among its neighbours by sending them a *ScanRequest* message. Upon reception of this message, the nearby APs switch to the MS's channel and listen to MS's packets for a certain length of time. APs that could receive packets from the MS reply to the current AP with a *ScanResponse* message including the RSS of the MS. The current AP chooses the AP with the highest signal and sends it a *StationMove* to designate it as the selected AP. Next, the current AP sends a message to the MS stating that it has to switch its channel to the channel of the new AP. The handover is transparent for the MS. The major limitation of this approach is that it assumes that AP are VAP; if not, the MS performs the costly IEEE 802.11legacy handover. It also assumes that all VAPs are connected to their neighbours, which makes this solution unrealistic for vehicular communications.

In general, these approaches suffer from a failure to take into account the motion of the MS in AP selection. For instance, the handover scheme described in [Khan 2012b] only requires the actual location of the MS while no information is required about its motion. It is reasonable to assume that AP selection will be different according to the direction of the MS.

2.2.3 Scanning inaccuracy

Although the impact of scanning on the duration of the handover can be significantly reduced and even eliminated, it still suffers from inaccurate results and partial information. There are multiple reasons for this. First, scanning is an instantaneous sensing of the network at one particular time, thus, it does not always reflect reality. Beacon loss can lead to the scan missing an available AP, especially in dense deployments, like those discussed in [Castignani 2012b, Arcia-Moret 2014]. Moreover, multiple empirical studies show that RSS changes appreciably from one probe to the next [Aguayo 2004]. In order to build a reliable picture of nearby APs' RSS it is necessary to take multiple samples and smooth the distribution. If the MS is moving, this method can also show RSS trends in the short term as described in [Mhatre 2006]. However, it requires the MS to spend a lot of time on each channel leading to long disconnections, or a focus on the current channel to the exclusion of others. In fact, the snapshot of the current network configuration provided by scans reflects reality only for a given period, which depends on MS mobility. In the case of a vehicle moving at 50 km/h the MS leaves the range of an AP in a few seconds. In addition, the MS can suddenly leave the LoS of the current AP (e.g. when the MS is hidden by a roadside artefact or a building, or when the vehicle turns). This makes scanning inefficient because it does not provide any information that could help to prevent such events.

In conclusion, as stated in [Mhatre 2006], the handover paradigm must shift from a negative definition (avoid imminent disconnection) to a positive one (selecting the AP that is most likely to provide good RSS in the near future). With this aim, other ways need to be found to provide information about the network environment to the MS. In addition, vehicular mobility characteristics such as trajectory limitations within defined areas (streets, car parks, etc.) can significantly ease predictions of location.

2.2.4 Predictive handover

In order to allow the MS to anticipate the short-term evolution of the network, it is critical to provide vehicle mobility information. Based on the fact that people usually drive on known routes, Deshpande et al. [Deshpande 2009] propose replacing scanning with AP selection using historical information. While driving, the MS learns and caches information about nearby APs (Extended Service Set ID (ESSID), BSSID, channel) by frequently probing networks during inactive periods. This information is used to script the handover location in advance such that the MS is always connected to the best candidate AP. Another approach is proposed by Kwak et al. [Kwak 2009], who suggest maintaining the vehicle's trajectory and neighbour information in order to select APs along the vehicle's path and predict an optimal handover location. This approach was investigated by Montavont et al. [Montavont 2006] who propose extending the Mobile IPv6 (MIPv6) architecture with a new component, a GPS server that is periodically updated with the location of the MS. Based on changes in the location of the MS and prior knowledge of AP locations, channels, BSSIDs and Internet Protocol version 6 (IPv6) prefixes, the GPS server triggers a handover to the closest AP when the MS is about to leave the range of the current AP by sending it a handover indication packet.

2.2.5 Reassociation optimization

2.2.5.1 IEEE 802.11f

The IEEE 802.11f trial-use standard, proposed in 2004 by the IEEE consortium recommends practices for multi-vendor AP interoperability. With this aim, IEEE 802.11f proposes the *Inter-Access Point Protocol (IAPP)*. The main objective is to maintain an association between an MS and a single AP during handover. At the first association, the MS sends an association request to the candidate AP. When the MS and the AP are associated, the AP broadcasts an *Add-Notify* message notifying the association with the MS. Upon reception of the *Add-Notify*, other APs clear any stale associations. When the MS hands-off to another AP, the MS sends a reassociation message containing the identity of the current AP to the candidate AP. The candidate AP sends a *Move-Notify* message to the current AP requesting the MS context and notifying the new association. The current AP replies with a *Move-Response* message.

Mishra et al. [Mishra 2004] observe that in the IEEE 802.11f context the transfer is performed reactively, which delays the association phase. In order to reduce this delay, the authors propose pro-actively sending the MS context to neighbouring APs. These are selected using NG, which capture the *reassociation relationship*, i.e. a potential handover from one AP to another, between access points. The NG is automatically generated by APs over time when they receive a reassociation request from an MS or a *Move-Notify* from another AP. The experiments conducted in this paper show that association latency is reduced from 15.37 ms to 1.7 ms. An improved version is proposed by Park et al. [Pack 2005] that attempts to reduce data overheads caused by sending the MS context to all neighbours. With this aim, the authors suggest evaluating the probability of possible transitions and sending the MS context only to those APs for which the association probability is higher than a given threshold.

2.2.5.2 IEEE 802.11r

IEEE 802.11r, otherwise known as Fast Basic Service Set Transition (FBSST) is applied when the network uses secure authentication. In IEEE 802.11 standards, during the handover, the MS reauthenticates and reassociates with the candidate AP. Thus, when IEEE 802.1x is used to secure communications, the MS is required to renegotiate its keys with the authentication server (e.g. a RADIUS server) at each handover. This key negotiation procedure significantly delays the handover. In order to reduce this, IEEE 802.11r allows the AP to cache the key from the server in order to avoid the full authentication process and only exchange the classic authentication request/ response messages. Bangolae et al. [Bangolae 2006] evaluate IEEE 802.11r in a basic setup of two APs. Their results show that the duration of the IEEE 802.11r handover is 42 ms versus 525 ms in a legacy implementation.

2.2.5.3 AP management centralization

The adoption of the Lightweight Access Point Protocol (LWAPP) by the IETF in February 2010 highlights the trend of centralizing IEEE 802.11 network management. Proposed by Airespace (subsequently purchased by Cisco), this protocol is available in all Cisco IEEE 802.11 networks. LWAPP proposes the transfer of multiple features from APs to an Access Controller (AC), which becomes responsible for bridging, forwarding, authentication, policy enforcement and optionally, encryption of user traffic on the network. The goal of LWAPP is to facilitate network management based on a single storage point for the network configuration (e.g. MAC lists, QoS allocation, AES authentication keys), decision-making and more efficient use of the computing power of APs.

During the handover, authentication and association requests are forwarded by the AP to the AC. Authentication/ association responses are sent by either the AP or the AC but the AC is always in charge of negotiating the key in a secure connection. The QoS provided to the MS is also negotiated by the AC. Therefore, unlike IEEE 802.11f-compliant networks, there is no need to transfer the MS context. To the best of our knowledge there is no publicly-available evaluation of the impact of LWAPP on the association/ reassociation delay.

2.2.6 Discussion

Vehicular communications rely on the potential of at least some vehicles to have Internet access. Currently, a vehicle can use either cellular or IEEE 802.11 networks that are densely deployed in urban areas. Although IEEE 802.11 networks provide high throughput and can respond to the speed of a vehicle, their range is short, leading to frequent and costly handovers between APs. This impact must be reduced in order to provide a continuous connection. Furthermore, as stated by [Mhatre 2006], the handover paradigm must change from a negative definition (avoid imminent disconnection) to a positive one (selecting an AP that is more likely to provide good RSS in the near future). This can be achieved by using other criteria to trigger the handover and optimizing the scanning phase. The scanning phase is the typical target of handover optimization; firstly because it is the most time-consuming, and secondly because it only provides a snapshot of nearby networks when mobile vehicles need data about their short-term evolution. In order to enable a predictive handover that can pro-actively select and associate with the best AP as soon as it become available (and not when the link with the current AP goes down), the MS must be provided with extra information about nearby networks.

Chapter 3

Enhancing the efficiency of IEEE 802.11 networks under vehicular mobility

---

Contents

3.1 Introduction

In current urban environments only cellular networks provide full Internet connectivity to an MS. However, IEEE 802.11 networks are expanding and becoming a credible alternative. According to statistics from South Korea, the United States, Canada, Japan, Germany, and the United Kingdom, the proportion of wireless traffic exchanged on Android smartphones through IEEE 802.11 networks has increased from 67% in August 2012 to about 73% in April 2013 [Media 2014]. However, multiple challenges remain to be overcome in order to seamlessly connect an MS to nearby IEEE 802.11 network, in particular if the MS is in a vehicle. As IEEE 802.11 networks have a short range, the MS must manage frequent handovers. These handovers need to be transparent for the mobile user (as is the case for cellular networks) to avoid disconnections. However, unlike IEEE 802.11 networks where handovers are triggered by the client based on locally collected information, in cellular networks handovers are triggered by the infrastructure (in particular an entity close to the base stataion), which uses multiple source of information to take its decision. For instance, in LTE networks the MS measures RSS downlinks on all available channels and sends them to the current cell through a *Measurement Report* message. The decision to handover is taken by the cell, based on RSS criteria and communication begins between the current and the candidate cell via the Mobility Management Entity (MME). Reception of the *Handover Response* means that the handover has been approved by the candidate cell and the current cell informs the MS that it has to hand-off to the candidate AP.

The centralized management of cellular networks significantly reduces handover latency at the MS. This is not the case for IEEE 802.11 networks where handover managed client side. In addition, in the context of vehicular communications, the AP selection process has to take into account a rapidly-changing network environment and prevent sudden disconnection due to the vehicle's motion. Client- and network-based handover techniques have been the subject of several contributions in the literature [Melia 2007, Kassar 2008]. However, in this thesis we only consider the client-based approach.

3.1.1 Context

Vehicular mobility characteristics Figure 3.1 compares the characteristics of three mobility patterns (indoor pedestrian, outdoor pedestrian and vehicular) and their effect on changes in RSS in order to highlight the specificities of vehicular mobility.

In terms of speed, there is a significant difference between vehicular mobility, which can exceed 100 km/h (although the average in urban environments is about 30 km/h) and pedestrian mobility that is in the order of 5 km/h. High speeds have a critical impact on RSS, which can dramatically change in a few seconds. Figure 3.2 shows changes in RSS for AP hotspots sensed by a vehicle moving at 5 m/s (18 km/h). In this figure, colors represent a given AP. The red line represents a threshold of -80 dBm below which link quality drops significantly. The circles highlight large variations in RSS over very short time periods. These variations correspond to an increase or decrease in the signal.

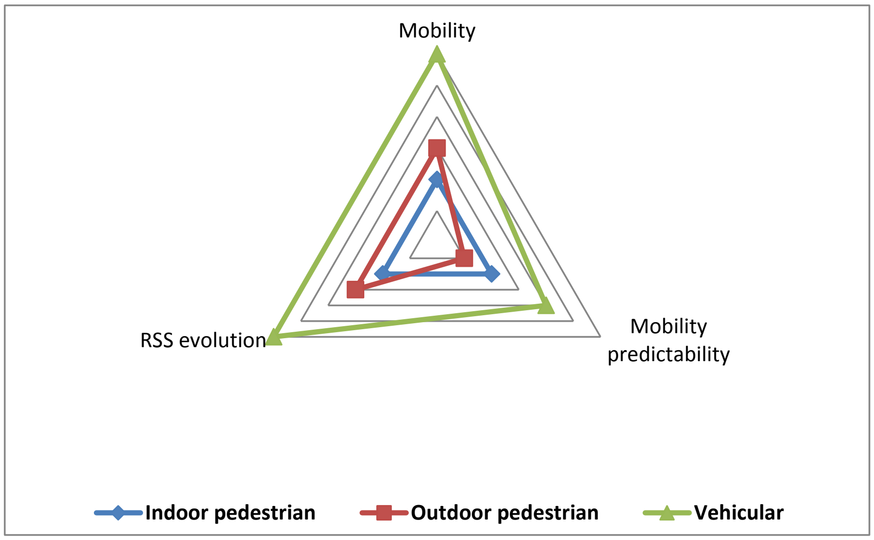

Figure 3.1: Comparison of mobility patterns

This makes the short-term prediction of RSS very difficult. However, changes in RSS are fairly smooth at lower speeds as the MS has enough time to analyze RSS trends, detect that the link is going down and find a suitable AP in order to perform a handover. This is particularly the case for indoor pedestrian mobility as the network has been designed to cover the whole space.

On the other hand, the pedestrian can move in any direction in two (sometimes three) dimensions. This makes pedestrian mobility, in particular outdoors, difficult to predict. This is not the case for vehicular mobility, which can often be reduced to one dimension(corresponding to the road infrastructure and consistent with driving regulations) and a choice between a given set of possibilities when the vehicle is at an intersection. As a result, vehicular mobility is much easier to predict than pedestrian mobility.

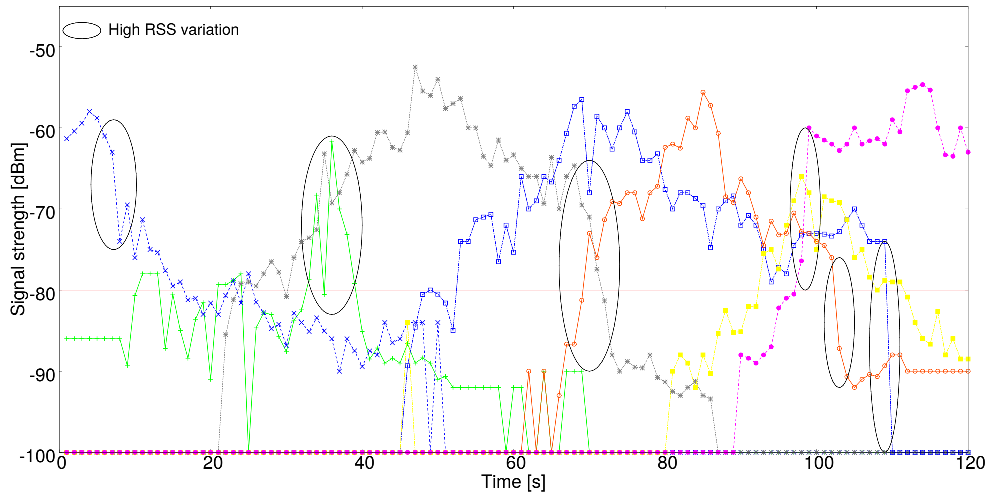

Figure 3.2: Changes in RSS under vehicular mobility

Legacy IEEE 802.11 handover under vehicular mobility We conducted a set of preliminary experiments to evaluate how standard IEEE 802.11 handovers in urban networks perform when used in vehicular mobility conditions. This preliminary study (i) provides a baseline for the legacy solution that can be compared with the solutions presented later in this section, (ii) evaluates the impact of vehicular mobility on MS communication, and (iii) evaluates the impact of the handover in the particular network considered in our work. Experiments were based on a commercial metro-scale IEEE 802.11 deployment (HotCity) in the city of Luxembourg. A detailed description of the HotCity network can be found in the next chapter. We used a WPA-supplicant (v0.73), which is a widely-used wireless connection daemon. As mentioned earlier, the trigger for the handover and AP selection are not defined in IEEE 802.11. In the case of a WPA-supplicant the handover is triggered when the RSS of the current AP drops below a given threshold. The WPA-supplicant selects the AP having the highest RSS as a candidate AP.

The following scenario was used: A vehicle moves at various constant speeds (15, 30 and 50 km/h) along an 875 m straight road. It accelerates, reaches its maximum speed and maintains this speed throughout the experiment. In order to mitigate the effects of other users at the same location accessing the network, we performed the experiment at night. The Internet connection was tested using an intensive TCP download from a web server.

Effect of speed and RSS handover threshold We conducted a comparative study of the duration of the connection and the data collected using various thresholds to trigger the handover. Figure 3.3 illustrates the mean throughput of the MS for different RSS thresholds (from −60 dBm to −85 dBm) at different speeds. This figure shows that vehicle speed has a critical impact on mean throughput. This is because the duration of the handover does not depend on the mobility of the MS when the time spent in the range of an AP is short. In addition, the trigger for the handover has a significant impact on received throughput. It appears that there is a trade-off between high RSS thresholds that produce too-frequent handovers, and low RSS thresholds that leave the MS connected to the current AP even when its RSS is too low to provide good link quality. As a result, we concluded that an RSS threshold of −75 dBm offers better average throughput along the path at all speeds.

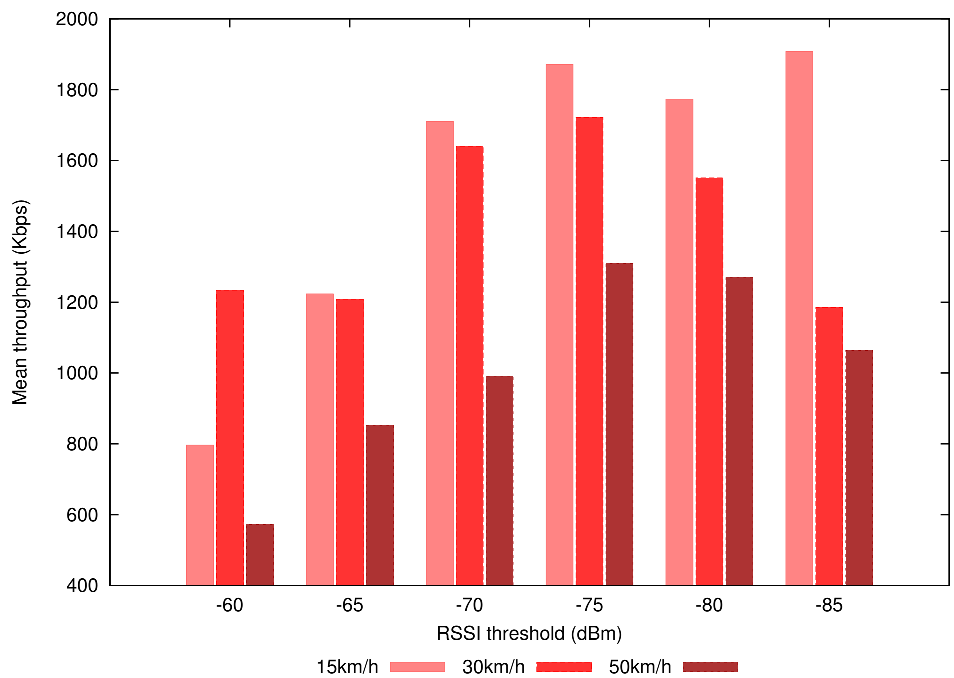

Figure 3.3: Mean throughput for different RSS handover thresholds

Handover duration Figure 3.4 shows the downstream TCP throughput of a representative run at 30 km/h. The color of the lines indicates the AP the MS is connected to. The vertical dashed lines illustrate instants when the vehicle passes closest to a given AP. Note that TCP throughput can be clustered into two distinct categories: high values (between 1500 and 3000 Kbps) and very low values (lower than 500 kbps). In all experiments, low values are observed almost exclusively when the handover is executed. After performing 15 experimental runs, we observed that an average of 3.6 handovers are executed during an average experiment lasting 105 s. The client remains connected to the same AP for 21 s and disconnected for 6.7 s on average. These results are consistent with related work [Deshpande 2010, Castignani 2012a]. Finally, the reconnection period was the main source of low throughput in our experiments.

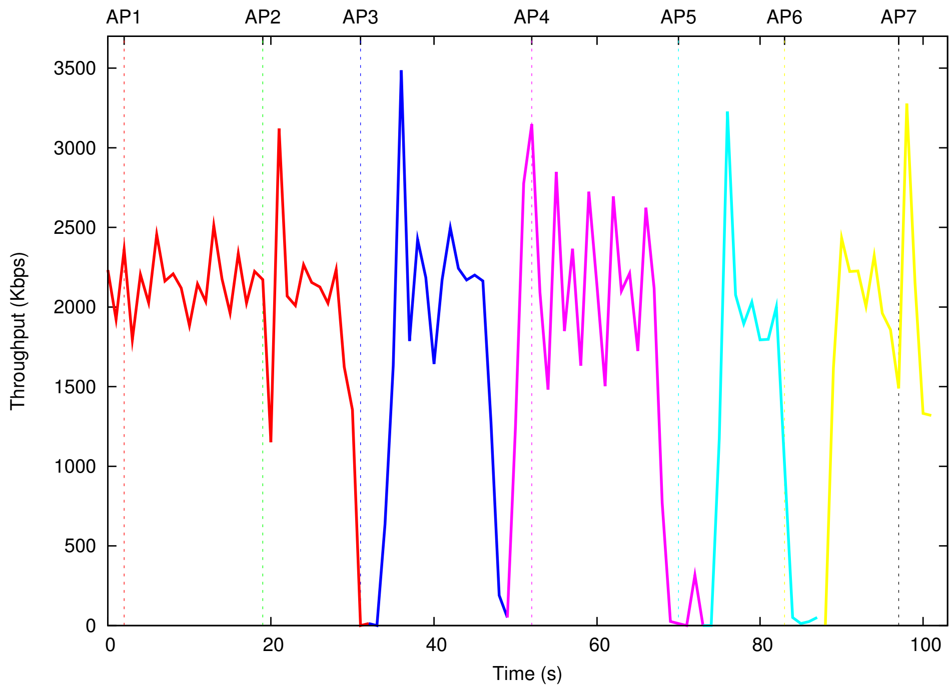

Figure 3.4: Downlink throughput

Our measurements show that handovers account for between 15 % and 30 % of the duration of the experiment. Moreover, we measured the contribution of the scanning process for each handover. This was calculated as the average time that elapsed between when the handover was triggered (i.e., current AP RSS lower than −75 dBm) and the beginning of the authentication phase, which here represents 75 % of the duration of the handover. This result is consistent with previous studies [Park 2004, Mishra 2004, Velayos 2004, Teng 2009, Castignani 2011a].

3.1.2 Objectives

These experiments highlight two ways to optimize vehicular communication in IEEE 802.11 networks. Firstly, a seamless MS handover and, secondly, the consideration of vehicular mobility during AP selection. In both cases, the scanning phase is the central element as it is the most time-consuming aspect of the handover, and because AP selection relies on the information provided by it.

Seamless handovers In order to optimize vehicular communications through IEEE 802.11 networks, the first objective is to eliminate the impact of the handover such that mobile users remain fully connected. This thesis focuses on the optimization of the Layer-2 (L2) handover. The most time-consuming phase of this handover is scanning, during which the MS periodically switches to all available channels to probe for nearby APs. Its duration depends on the number of channels and the time spent scanning each channel. The IEEE 802.11g standard specifies up to 14 channels. For IEEE 802.11n, which operates in both 2.4GHz and 5GHz channels, this number raises to 30 [FCC 2014]. As a result, the legacy scanning procedure leads to long delays that are not compatible with the objective of seamless transitions. Eliminating its impact is, thus, a priority.

Efficient AP selection In current urban environments, the density of IEEE 802.11 APs is high enough for the MS to frequently be within the range of multiple APs [Mota 2013, Castignani 2013]. The MS needs to find the one that performs best based on the always best connected paradigm [Gustafsson 2003]. This task is difficult to achieve if we only consider scanning information. The RSS snapshot of nearby APs is only valid for a very short time if the vehicle is moving. In order to select the best AP, i.e., the AP providing the highest RSS in the near future, it is necessary to have information about short-term changes in RSS of nearby APs. Based on this information, the MS should be able to select a candidate AP as soon as this AP becomes the best option. This leads to a change in the handover paradigm. Rather than triggering the handover when the MS considers that it is necessary due to an imminent disconnection, it occurs when a new best AP is detected. The aim of this approach is to optimally exploit nearby IEEE 802.11 networks by maximizing the RSS of the MS.

3.1.3 Propositions

In this thesis we make three proposals intended to make the handover seamless and to optimize the AP selection process.

A second interface for seamless handovers The introduction of a second interface has been proposed before in the literature. Some of this work is described in Section 2.2.2.5. The approach makes it possible to achieve the first objective of this work i.e., remove the impact of the handover by handling the scanning and authentication/ association in the second interface while the first remains connected to the network. However, this does not enhance AP selection policy. Therefore, although it is a suitable approach for vehicular communication, it does not alone fulfil all the requirements. The implementation of the approach is used to evaluate on the one hand, gains in terms of connection time, and on the other hand, the negative impact of scanning-based AP selection.

Replacing scanning by prior knowledge of network characteristics As mentioned earlier, the scanning process slows down the handover significantly and does not provide enough information to allow the MS to choose the best AP. To overcome these issues, we propose removing scanning from the handover process. It is replaced by prior knowledge of the network topology and a simulated RSS. Based on this information, the MS can deduce from its current location the most suitable AP in all possible directions. In out work, this happens when a vehicle approaches an intersection and the AP chooses the best APs in the available directions.

Including vehicular mobility predictions in AP selection In most cases, only the driver of the vehicle knows the destination and the route that will be taken. This is an issue for AP selection since, even if the MS has prior knowledge of the location and RSS of all of the APs along the entire route, the best AP depends on the direction taken at an intersection. In order to provide optimal AP selection, there is a need to predict this direction. This must be performed on the fly since the MS has no information about the vehicle's route. The method proposed in this thesis is based on an analysis of the vehicle's trajectory. Knowledge acquired on the fly allows the MS to choose the best possible APs, based on prior knowledge of network topology and a simulated RSS, which corresponds to the direction taken by the vehicle.

3.2 Introducing a second interface: evaluation of a make-before-break approach in vehicular communication

An MS using the standard IEEE 802.11 handover process needs to perform several actions outside the channel of the current AP in order to be granted access to the candidate AP. Although they are necessary, there is no requirement for them to be performed by the interface connected with the current AP. Therefore, it is possible to relocate them to a second interface. As there are no space or energy constraints in vehicles, a second radio may be a suitable solution for vehicular communications. In this case, the handover to the candidate takes place before the link is broken with the current AP; this is called the *make-before-break* approach.

3.2.1 The make-before-break approach

The solution considered in this thesis is based on the algorithms proposed by Ramachandran et al. [Ramachandran 2006]. The authors propose a make-before-break mechanism relying on a second IEEE 802.11 radio. Two L2 roaming algorithms are described: a two-card static algorithm (referred to here as static, indicating the "static roles" algorithm), and a two-card dynamic algorithm (referred to here as dynamic, indicating the "dynamic roles" algorithm).

The static algorithm The static algorithm completely decouples scanning and the data-connection roles and creates separate interfaces (namely, the *scanning interface* and the *data-connection interface* respectively). As depicted in Figure 3.5(a), the *scanning interface* continuously probes all channels, gathers information about nearby APs and decides whether one of the candidate APs is better than the current AP in terms of RSS. We refer to this candidate AP as the *best AP*. The *scanning interface* shares information about the best available APs with the *data-connection interface* through a common resource, a description file located in the user space. The *data-connection interface* first tries to associate to the best AP. If the connection attempt fails, it rechecks the common resource for a new best AP and attempts a new connection. On the other hand, if the connection succeeds, the MS performs data communication through the new AP. The *data-connection interface* continues to make periodic checks of the best AP information on the common resource. When the best AP changes, it tries to roam to it. The message exchange during the handover using the static algorithm is depicted in Figure 3.6.

During the handover, the *data-connection interface* simply performs authentication and association with the best AP discovered by the *scanning interface*. This reduces the length of the handover to authentication and association latency. It also provides better knowledge of nearby APs, since the *scanning interface* can retrieve AP information more often without caring about the impact of scanning. However, as the algorithm is based on a break-before-make handover approach the MS can still be disconnected, even if for a much shorter time.

The dynamic algorithm The dynamic algorithm aims to completely eliminate disconnection during handover (see Figure 3.5(b)). For the first connection, the *scanning interface* provides the best AP information to the common resource. Then, the *data-connection interface* attempts an association with the best AP. If this connection succeeds, the *data-connection interface* keeps exchanging data until the *scanning interface* performs a new connection to a different AP. This occurs because the *scanning interface* will continuously look for a new best AP and attempt the connection by itself, without asking the *data-connection interface* to attempt the new connection. Note that at this particular moment, the interfaces switch roles. If the *scanning interface* fails in attempting the new connection, it will first perform a new scan before trying again to associate. When the roles of the interfaces are exchanged, Layer-3 (L3) settings (IP addressing and the routing table) are transferred from the former *data-connection interface* to the new one. The message exchange during the handover using the dynamic algorithm is depicted in Figure 3.7.

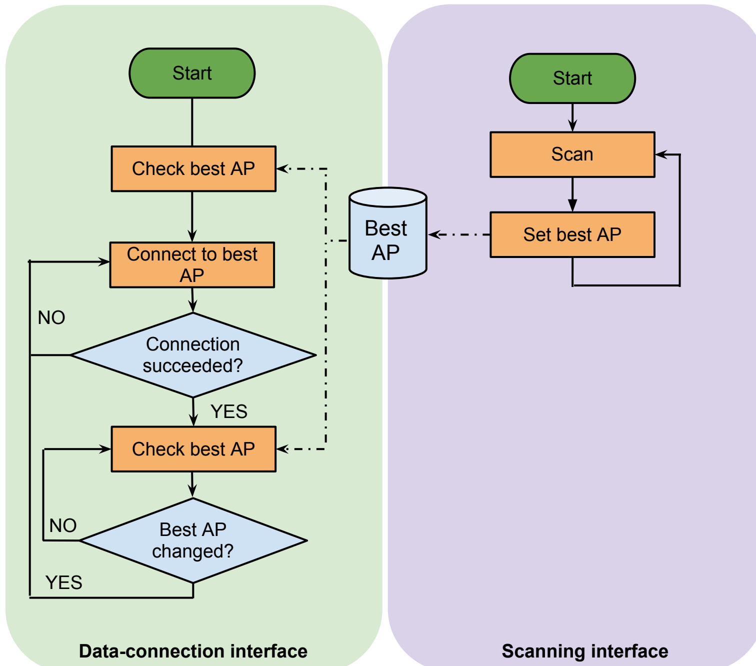

(a) Static algorithm

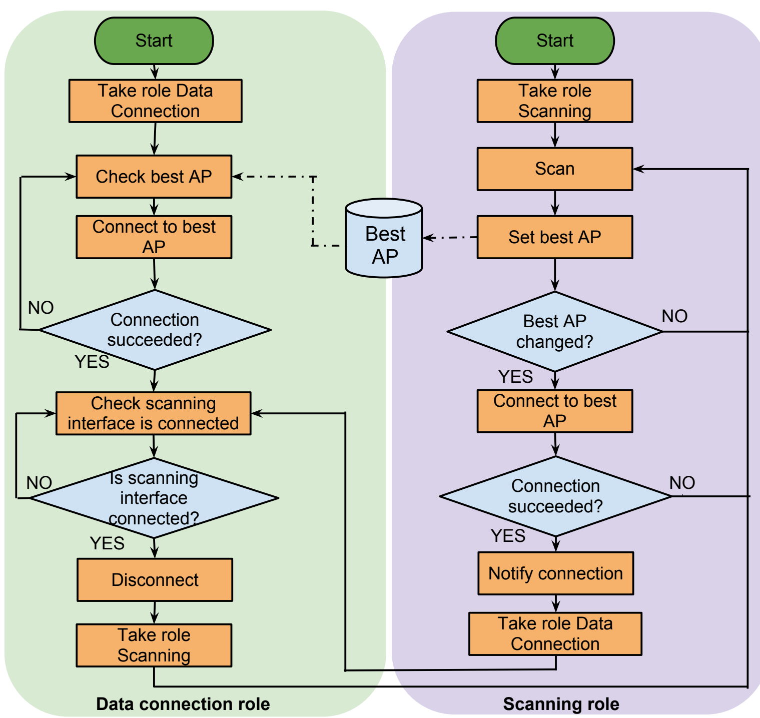

(b) Dynamic algorithm

Figure 3.5: Implementation of static and dynamic algorithms

Unlike the static algorithm, if there is a good candidate AP to roam to, the MS will no longer be disconnected, and the handover has no effect on client connectivity. However, the algorithm implies that the MS is simultaneously connected to the same network through two different interfaces. This issue requires important changes either on the client or the network side. The *scanning interface* can connect to the network at the same time as the *data-connection interface* if and only if one of the two following conditions is met. First, the network allows two simultaneous connections with two interfaces from the same client and allocates them the same IP address. Second, the client hides the fact that it has two interfaces by assigning them the same MAC address. As we used a commercial network for our experiments, we were not allowed to introduce the first modification. Thus, we decided to use the same MAC address for both interfaces.

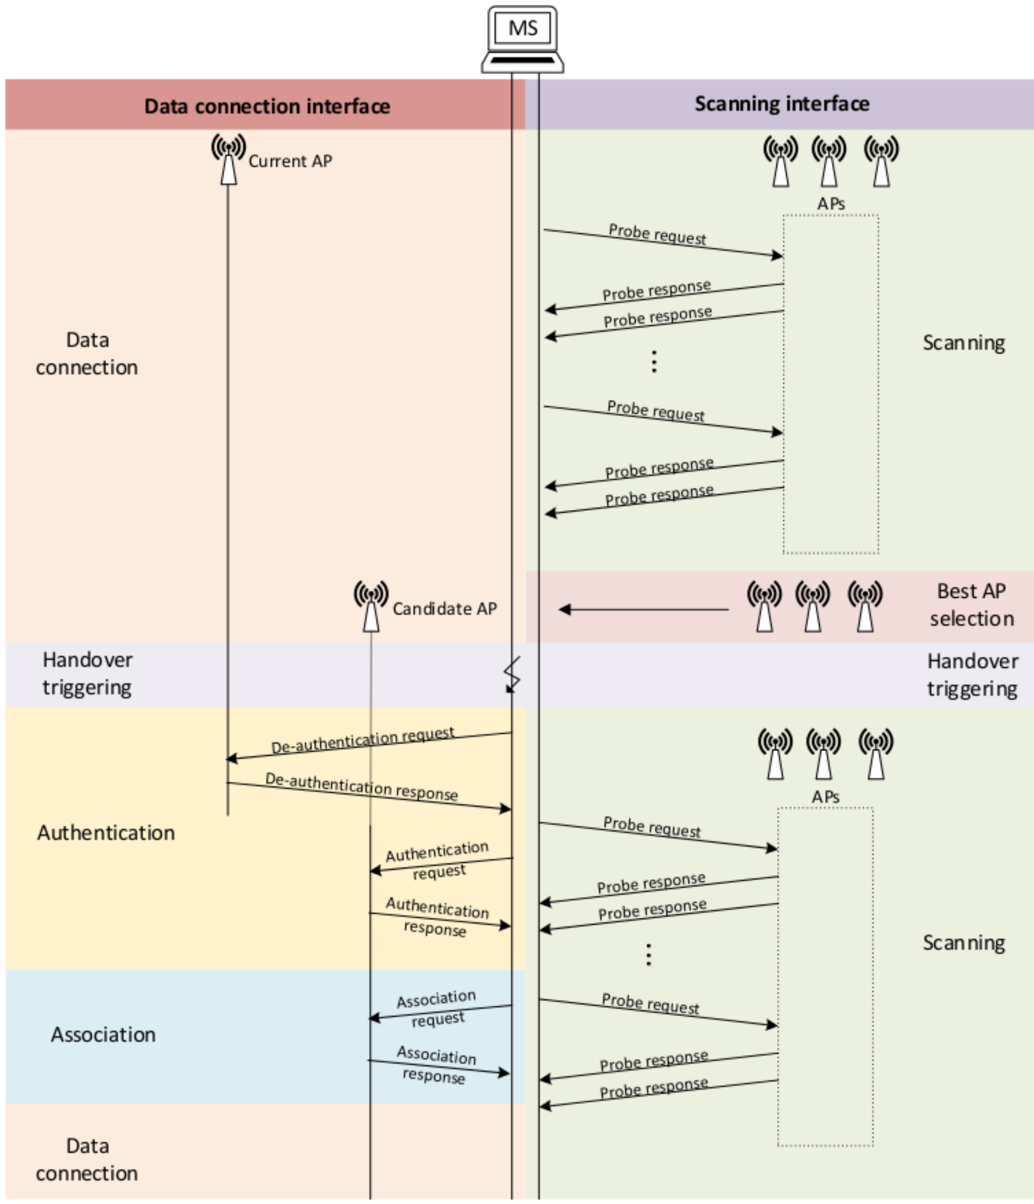

Figure 3.6: Handover using the static algorithm

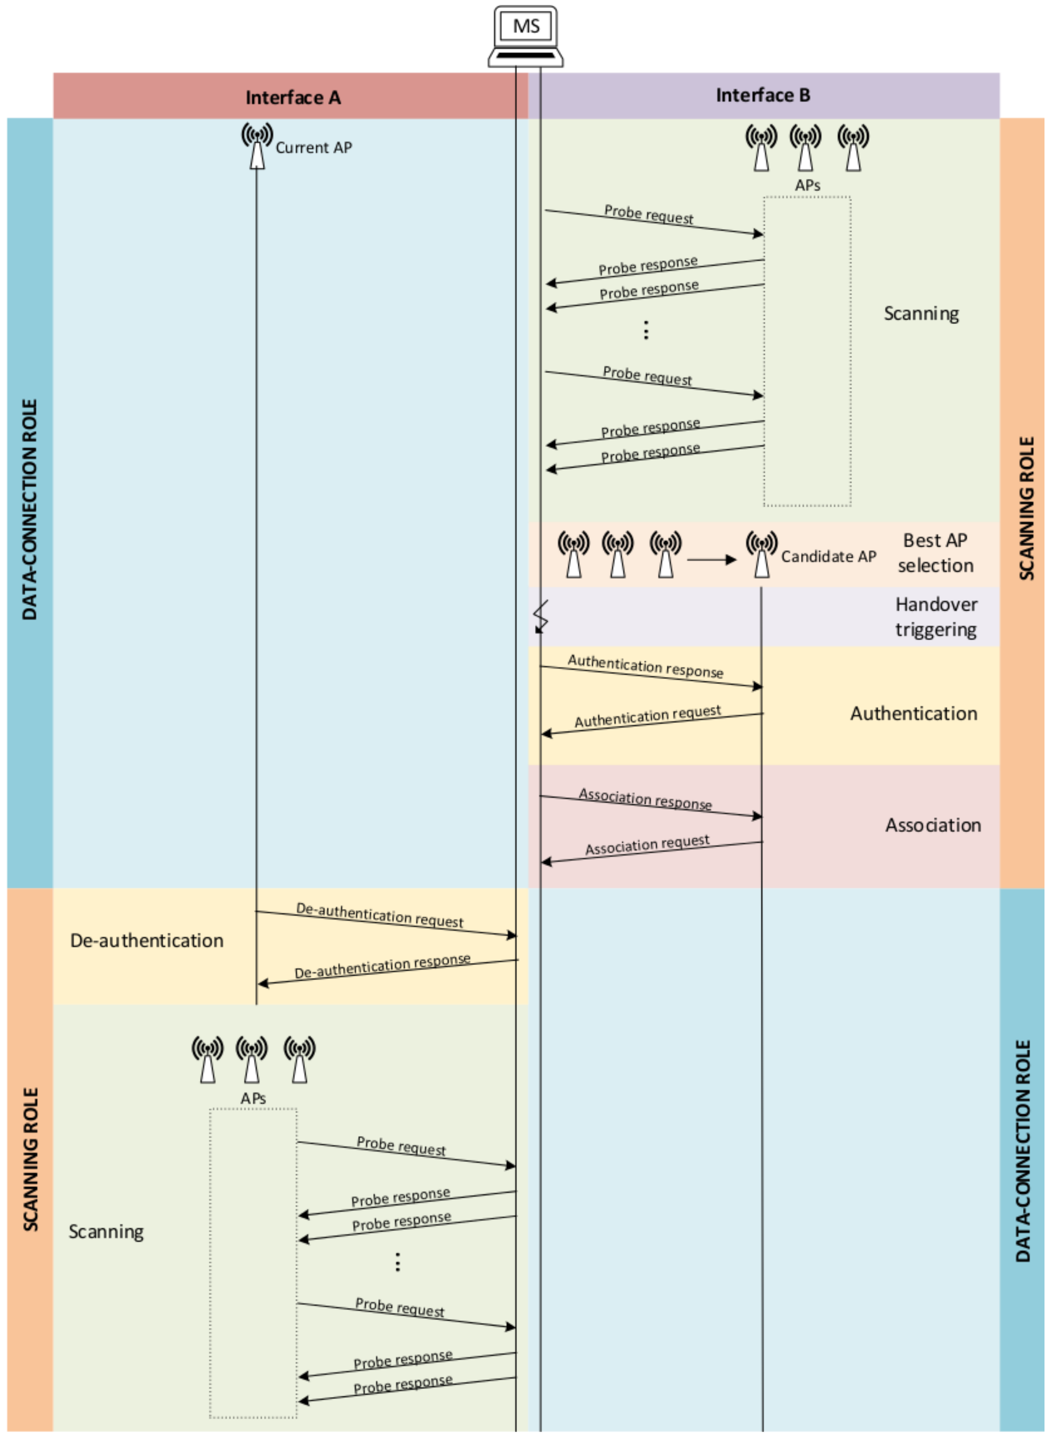

Figure 3.7: Handover using the dynamic algorithm

3.2.2 Implementation

We implemented these two algorithms in the WPA-supplicant v0.73, a widely-used wireless connection daemon. Experiments were performed on laptops using Linux Kernel v2.6.35. Note that as they are implemented in the WPA-supplicant and not in the driver, the algorithms can run on different wireless cards with different chipsets. We extended the WPA-supplicant to manage the connection to the network with two interfaces. Then, we replaced the main thread of the legacy version with a two-thread implementation: one managing the scanning (the *scanning thread*) and the other managing the connection (the *data-connection thread*). Note that we use the existing scanning, authentication and association primitives and we invoke them in our two-thread implementation. The scanning thread invokes the scanning primitive, and then processes and stores the information gathered from the different APs (BSSID, channel, RSS). Depending on the results, it chooses a candidate AP and shares it with the data-connection thread. The data-connection thread obtains information about the selected AP from the scanning thread and attempts to connect to it. We used a basic rule to manage roaming: when the RSS of the current AP falls below −75 dBm the data-connection thread looks for a candidate AP with a higher RSS.

3.2.3 Scanning inaccuracy under vehicular mobility

This section highlights partial results related to scanning efficiency. The path used in these experiments was 1.8 km long and traversed 15 APs. These APs offer the best QoS as the MS can stay in their LoS for long periods. They are labelled from 1 to 15 according to their order of appearance. There are also multiple other APs that operate in the streets adjacent to the experimental path. We tested three handover approaches: the Single Interface (SI) approach, which is the legacy solution presented in Section 3.1.1, and the static and dynamic algorithms using a second interface.

First, we consider the SI approach. Table 3.1 shows the sequences of AP chosen by the MS based on scanning. Their length varies from 10 to 13 APs. Note that they are all different, which shows the instability of scanning results. We conclude that the QoS provided by the AP using this approach on the same route may vary from one run to another.

Next we consider scanning results using static and dynamic algorithms. The scanning thread is set to scan very frequently (more than twice per second) based on a short Maximum Channel Time (MaxCT), i.e., 25 ms. Unlike the SI approach, no threshold was defined to trigger the handover. This method is intended to enhance AP selection by increasing the amount of information provided by the scan and trigger the handover as soon as the candidate AP becomes the best AP. However, we observed that the best AP differs from previous scan results in more than 31% of cases. Of these, 20% of APs are not located on the vehicle's route. This confirms the instability and the inaccuracy

Experiment #	Sequences of APs
1	1 3 5 6 7 8 10 11 13 14 15
2	1 2 3 4 6 8 10 12 13 14 15
3	1 2 3 4 6 7 8 10 11 13 15
4	1 3 4 5 6 7 9 10 12 13 14 15
5	1 3 4 5 6 8 10 11 12 13 14 15
6	1 3 5 6 7 8 10 12 13 14 15
7	1 3 4 5 6 7 8 11 0 13 14 15
8	1 3 4 5 6 7 8 10 11 12 13 14 15
9	1 3 4 6 7 8 11 12 13 14 15
10	1 3 5 6 7 10 11 12 13 14 15

Table 3.1: Sequence of APs in ten experiments using a single interface

of scanning shown in SI results. In addition, this phenomenon implies an increase in the number of handovers (3.5 times higher compared to the SI approach). Note that this number could be partially reduced by introducing regulation mechanisms such as a hysteresis process.

In conclusion, scanning instability means that AP selection cannot be based on raw results. Although filtering would rule out some APs that do not provide a reliably high RSS, the approach does not prevent sudden variations in RSS, as depicted in Figure 3.2.

3.3 Replacing scanning with prior knowledge of the network

3.3.1 Basic principles

Vehicular communications suffer from frequent disconnections when using IEEE 802.11 networks. In the literature, several solutions have shown how to significantly reduce the impact of the handover in terms of disconnection time, like the make-before-break approach presented in the previous section. However, these solutions, when they rely on scanning, do not address the issue of inefficient AP selection under vehicular mobility. In order to reduce the impact of the handover and improve the AP selection process, we propose MROAD, a Model-based ROaming Decision technique. MROAD is intended to provide always-best connected wireless access without scanning during the handover. Instead, it uses prior knowledge of IEEE 802.11 AP RSS along the vehicle's route, obtained through simulation. The approach provides an optimal handover sequence to the MS. Specifically, with this new source of information, the MS is able to select the AP with the highest RSS at any place along the route. This solution allows the handover

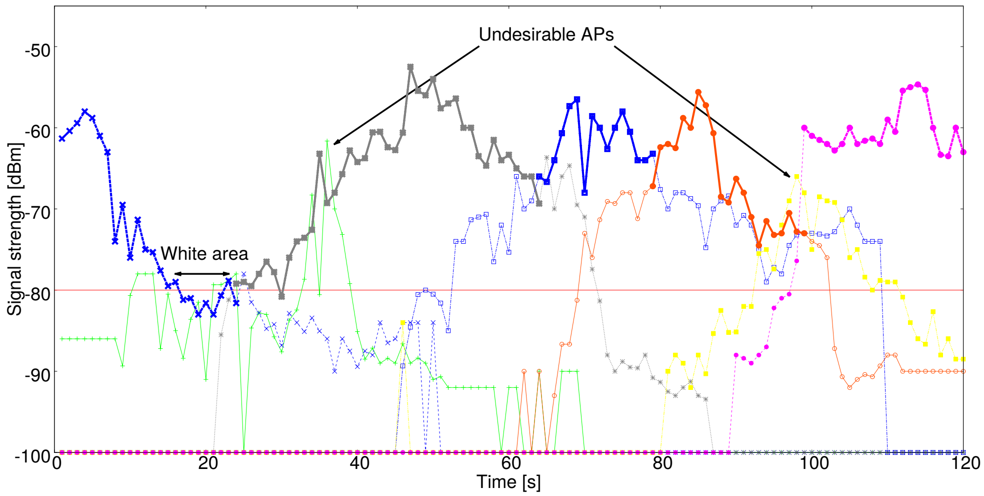

Figure 3.8: Changes in RSS under vehicular mobility

to be triggered at the optimal location, i.e., where the candidate AP becomes the best AP. Figure 3.8 show the changes in RSS in a hotspot network when the vehicle moves at approximatively 5 m/s (18 km/h). Prior knowledge of changes in RSS enables: (i) the detection of undesirable APs, i.e. APs the do not have a reliable RSS, (ii) selection of the best AP (those shown with wider lines in Figure 3.8) and, (iii) anticipation of white areas, i.e. areas with excessive packet loss due to low RSS.

Although MROAD should significantly reduce the length of the handover, it may have an impact on MS communication. In addition, triggering the handover when a candidate AP becomes the best AP may increase the frequency of handovers. This is unlike legacy solutions that wait until the RSS of the current AP becomes so low that an alternative has to be found. Therefore, on one hand the proposed method may have a positive impact: increased RSS, but on the other hand a potential negative impact: longer disconnections due to more handovers. This trade-off must be taken into account during the AP selection process in order to optimize MS communications. The solution proposed here consists of pre-computing the AP selection offline based on simulations.

3.3.2 Building prior knowledge of the network

There are many ways to characterize IEEE 802.11 networks. The most accurate consists of sensing the entire area they cover. A collaborative approach is proposed in [Deshpande 2009]. In such an approach, the MS periodically senses the network when it is idle. It has the advantage of providing accurate information dynamically, but requires a large community in order that information is frequently updated and remains accurate. In contrast, our approach consists of modeling the RSS through a simulation platform. However, it has multiple drawbacks. First, it requires knowing the location and characteristics of APs. Second it offers a static view of the network that may become obsolete if an AP is down or if its configuration is dynamic (e.g. due to automatic channel allocation). On the other hand, it has the advantage of simplicity. It can quickly provide results with no need for a community that must frequently sense the network, and it is scalable. The accuracy of the RSS model depends on the accuracy of the information and the model used. The method used in our research is shown in Figure 3.9 and described hereafter. First, the map of the experimental route is extracted from OpenStreetMap [1]. OpenStreetMap is a collaborative project that aims to map the world. Its web interface provides links for downloading a detailed map of an area selected by the user. Beyond a description of the road network, OpenStreetMap provides information about all other transport infrastructure (e.g., train, bus, bicycles), a description of the urban topology (e.g., buildings, parks, pedestrian areas), and points of interest (PoI). In addition, we obtained a list of APs and their locations from an open database provided by the local administration [2]. The mobility scenario is then simulated with SUMO [DLR], a urban mobility simulator. The OpenStreetMap map is converted into SUMO files by the *netconvert* tool provided in the SUMO suite of tools. Finally, the scenario is exported to Qualnet [Networks], a network simulator, using the Vergilius framework [Giordano 2010a]. We use Qualnet to simulate the RSS of APs along the vehicle's route. RSS is computed using CORNER [Giordano 2010b], an urban propagation model based on knowledge of the road topology. CORNER computes signal attenuation taking into account three situations: (i) the AP is in the LoS of the MS, (ii) the AP is separated from the LoS of the MS by a corner, or (iii) the AP is separated from the LoS of the MS by two corners. The model provides a good trade-off between accuracy and computational cost. The MS splits the vehicle's route into small segments and uses the simulated RSS to compute the average signal quality on each segment. In our experiments, the size of the road segment is set to 5 m, which is the highest possible accuracy from a GPS location update. A model of the RSS for a HotCity AP is shown as an example in Figure 3.10. The MS then runs the AP selection algorithm described in the following section.

[1] http://openstreetmap.org
[2] http://www.topographie.lu/

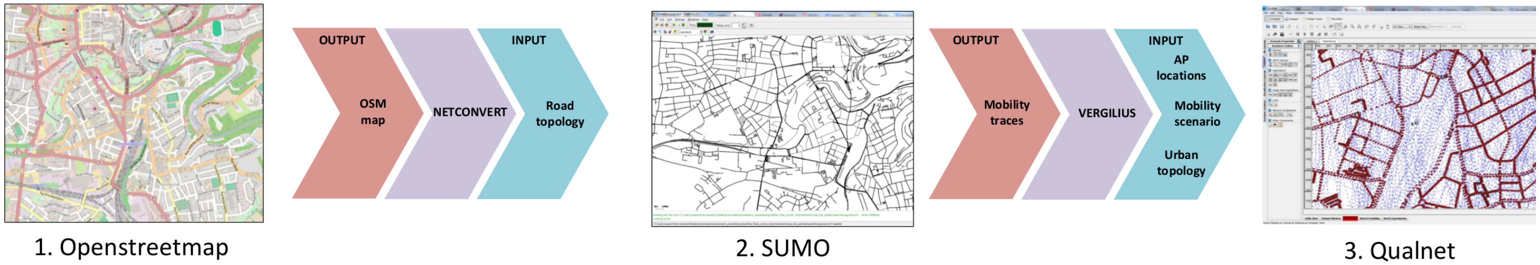

Figure 3.9: The three stages of modeling the RSS of IEEE 802.11 APs

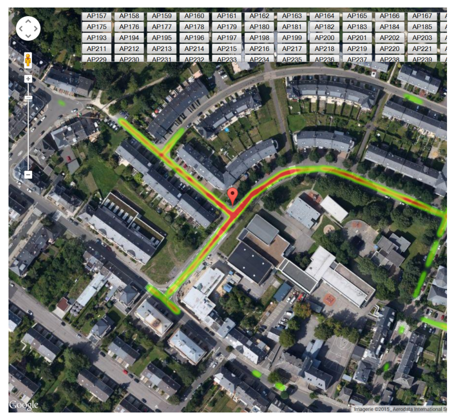

Figure 3.10: RSS modeling of a HotCity AP

3.3.3 MROAD AP selection

MROAD is based on two assumptions: (i) the MS knows the vehicle's route and the location of APs in advance; and (ii) the vehicle does not pass the same place twice. The first assumption makes it possible to pre-compute roaming locations but implies prior knowledge of the vehicle's destination. The second assumption was included to simplify the implementation and evaluation of the solution and could be removed without impacting the principle of the approach. AP selection is performed in two steps. First, MROAD builds an overlapping AP graph using the simulated AP RSSs. This overlapping graph provides all of the possible sequences of feasible candidate APs for a given scenario. For each sequence, MROAD uses the graph to find the best handover location, i.e. where a candidate AP provides a better RSS than the current one. Then, MROAD selects the best AP sequence from all of those available (where the time the MS has Internet access under good conditions is highest). When the vehicle is on the road, the handover is triggered based on its location.

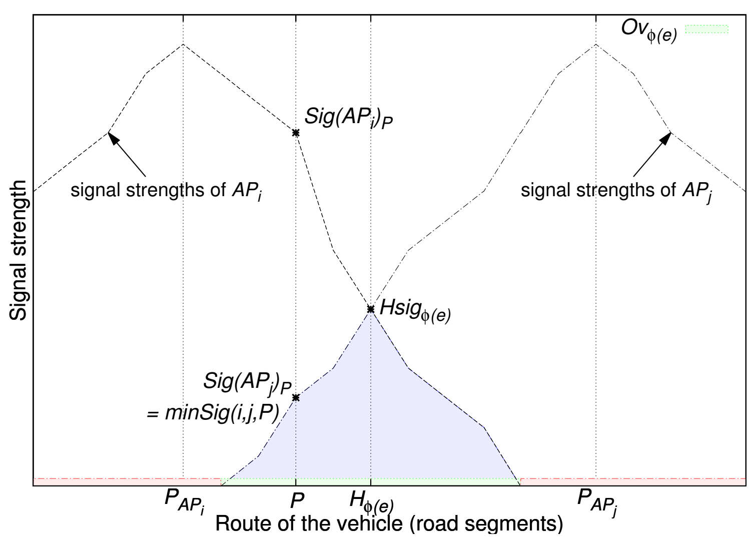

Figure 3.11: Handover location computation for AP_i and AP_j

Overlapping AP graph MROAD creates an overlapping AP graph based on predicted signals in order to compute sequences of APs that are able to provide IEEE 802.11 connectivity along the vehicle's route. Let the 3-tuple $G = (V, E, \phi(E))$ be an overlapping AP graph based on simulated AP RSS where:

- V is the set of vertices of G, i.e., the set of APs on the vehicle's path.

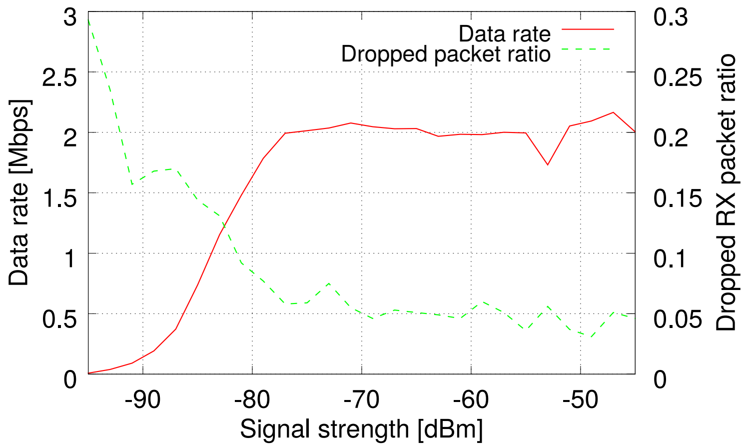

Figure 3.12: Relationship between RSS, data rate and dropped packets

- E is the set of edges of G, i.e., the set of connections in the graph between two overlapping APs on the vehicle's path.
- ϕ is the *incidence function* of G

Let $V = \{AP_1, AP_2, ..., AP_n\}$ be the vertex sequence of all APs in the vehicle's path sorted by order of appearance, P_{AP_i} is the road segment where AP_i is located, $Sig(AP_i)_P$ is the RSS of AP_i in the road segment P and $minSens$ is the network interface sensitivity. The two elements of $\phi(e) = \{AP_i, AP_j\}$, for any $e \in E$, are the vertices of the edge e, i.e., the overlapping APs linked in the graph by the edge e. Figure 3.11 is an example of the received RSS for the transition from AP_i to AP_j. In this figure, we can see P_{AP_i}, P_{AP_j} and for the road segment P, $Sig(AP_i)_P$ and $Sig(AP_j)_P$. MROAD considers that there is an edge from AP_i to AP_j in the overlapping AP graph if there is connectivity continuity from AP_i to AP_j. As a result,
for $i, j \in [1, n], i < j, e = \phi^{-1}(\{AP_i, AP_j\}) \in E\ if$

$$\exists\ P\ where \begin{cases} Sig(AP_i)_P > minSens \\ Sig(AP_j)_P > minSens \end{cases} \tag{3.1}$$

Eq. 3.1 requires that the range of the APs linked in the graph overlap, i.e., at a given place, the MS is able to detect signals from both APs. As stated above, the vehicle does not pass the same place twice. Therefore, in order to simplify the AP selection process, we define the overlapping AP graph as an oriented graph such that the connection from AP_i to AP_j is possible only if AP_j follows AP_i in order of appearance.

Handover location computation We define the handover location as the road segment $H_{\phi(e)}$ that provides the best conditions in terms of RSS to perform a handover.

In this work, we consider that the best conditions are met when the RSS of both APs are highest levels. This choice was based on the observations in [Aguayo 2004, Mhatre 2006, Castignani 2012a] and confirmed by our experiments. Figure 3.12 shows that there is a correlation between RSS and the MS downlink. Similarly, Figure 3.12 shows that packet loss decreases as RSS increases.

As a result, although the candidate AP is known and may provide satisfactory RSS from a given moment, it is optimal to perform the handover only when the current AP provides higher RSS. This allows the MS to benefit from the best available RSS and the association with the candidate AP to be performed under the best conditions, i.e. when packet loss is lowest. This approach should mitigate impacts on MS throughput.

Giving the preceding arguments, we force handovers to occur in locations where the RSS of both the current and candidate APs are highest. Note that the word location refers here to Road Segment (RS) of around 5 m defined during the simulations. The handover location is computed offline, based on simulated RSSs. It consists of three steps. The first is the definition of the overlapping area, i.e., the area where the MS is able to receive a signal from both potential APs. Let $Ov_{\phi(e)}$ be the overlapping area of the set of road segments P that satisfy Eq 3.1. In Figure 3.11, the overlapping area is illustrated by a green line. The second step is the selection of the lower RSS for each road segment in the overlapping area from the potential APs. This is performed on each road segment in the overlapping area. Let $minSig(i, j, P)$ be the lowest signal function such that for $P \in Ov_{\phi(e)}$

$$minSig(i, j, P) = min(Sig(AP_i)_P, Sig(AP_j)_P) \tag{3.2}$$

In Figure 3.11, for the road segment P, $minSig(i, j, P)$ is $Sig(AP_j)_P$. The set of all the elements with the lowest signal function is shown by the curve surrounding the blue area. The final step is the selection of the handover location as the road segment that maximizes the lowest signal function. Let $H_{\phi(e)} \in Ov_{\phi(e)}$ be the handover location, i.e., the location in the overlapping area in which the handover signal quality $Hsig_{\phi(e)}$, is:

$$Hsig_{\phi(e)} = max_{P \ \in \ Ov_{\phi(e)}}(minSig(i, j, P)) \tag{3.3}$$

Figure 3.11 shows the handover signal quality $Hsig_{\phi(e)}$ as the maximum of the curve surrounding the blue area and the handover location $H_{\phi(e)}$, which is the road segment corresponding to $Hsig_{\phi(e)}$.

Computation of optimal sequences The AP selection graph provides the MS with a list of APs with continuous connectivity. Therefore, starting from the first element of the graph, the MS has several possible transition sequences that lead to the final AP. We evaluated these sequences in order to select those with best connectivity. This evaluation took into account the handover trade-off discussed above: namely increased RSS and longer disconnections. Therefore, we evaluated sequences based on two parameters. The first is the signal quality of each handover. This parameter is a good indicator of overall RSS as it is a local minimum that takes into account the RSS of successive APs. The second parameter is the length of the sequence, i.e. the number of handovers.

The aim was to compute a score for all handover sequences that combined these two parameters. This score is a bi-objective optimization problem, described below. Let $E_N \subseteq E$ be a handover sequence, where N is the index of the sequence. We evaluate this sequence by calculating $score(E_N)$. Let $score(E_N)$ be the sequence score of E_N, $|E_N|$ the sequence length of E_N and $Trans(E_N)$ a metric of the quality of transitions between the different elements of E_N.

$$score(E_N) = \{|E_N|; Trans(E_N)\}$$

We define $Trans(E_N)$ as the average handover signal quality of transitions in the sequence:

$$Trans(E_N) = \sum_{e \in E_N} \frac{Hsig_{\phi(e)}}{|E_N|} \tag{3.4}$$

These two parameters highlight two, opposing trends. When sequence length is the only consideration, an optimal sequence contains APs that are as far from each other as possible in order to decrease the number of handovers. However, this may have an impact on RSS and performance (e.g., low throughput and multiple association attempts due to frame loss). On the other hand, when APs are closer to each other, the quality of the handover improves but disconnections are more frequent. Consequently, the trade-off between the number of handovers and signal quality must be managed. We introduce $Trans(E_N)$ to handle this. The bi-objective function, $score(E_N)$ is optimized by minimizing $|E_N|$ and maximizing $Trans(E_N)$. In order to reformulate the optimization problem into a pure minimization problem, we apply the duality principle [Deb 2001] and consider $-Trans(E_N)$. Now, the bi-objective problem optimizing $score(E_N)$ consists of minimizing $|E_N|$ and $-Trans(E_N)$. Figure 3.13 shows the set of feasible solutions for this problem. In particular, we consider the Pareto-optimal front, which is composed of seven handover sequences (from E09 to E15), with different values for $score(E_N)$. The experiment compared the performance of each Pareto-optimal front sequence. First, it focused on the relevance of the parameters involved in the computation of $score(E_N)$. It then evaluated the trade-off between minimizing the number of handovers and maximizing their quality.

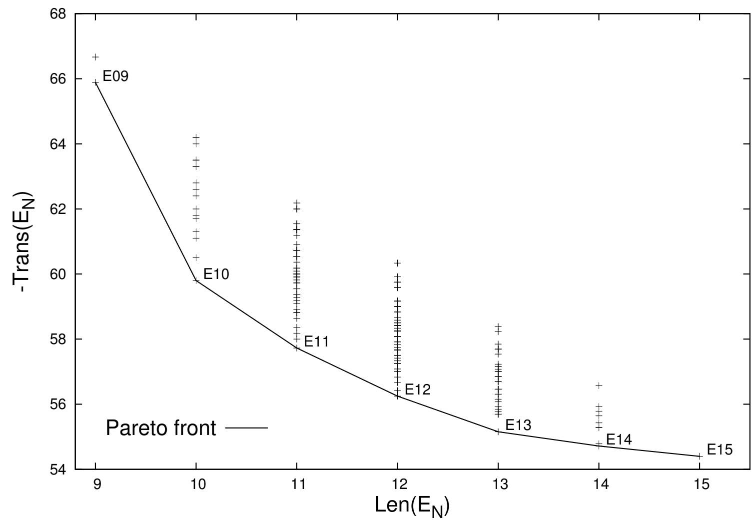

Figure 3.13: Distribution of $score(E_N)$

3.3.4 Implementation

In order to evaluate the handover sequences, we implemented connection software. Our implementation is based on the WPA-supplicant architecture and uses Linux Wireless EXTensions (WEXT)[3] as an interface to the kernel space. It is composed of a set of primitives responsible for probing, connecting and performing the roaming decision. The roaming decision primitive is the main component of MROAD. It takes as input a list of road segments for the route in question, the handover sequence with corresponding handover locations, and the current state of the MS (location and connection status). The roaming decision primitive processes these inputs by selecting the AP corresponding to the location of the vehicle and decides if a handover needs to be executed.

Its output is connection primitive information used to connect to the selected AP. The location of the vehicle is provided by an external USB GPS receiver placed inside it, close to the windscreen. This receiver provides a location update with an accuracy of 5 to 10 m every second. As the vehicle can reach speeds of up to 50 km/h, it can travel up to 13.9 m in 1 s. Consequently, it is necessary to update the vehicle's location more frequently in order to obtain more accurate location information. To this end, we implemented a prediction mechanism that approximated the vehicle's location when the roaming decision primitive has no recent GPS fix. This simple mechanism takes the vehicle's speed and direction over the past second and computes the current location by multiplying the elapsed time since the last GPS fix by the vehicle's speed over the

[3] http://wireless.kernel.org/en/developers/Documentation/Wireless-Extensions

past second. We assume that vehicle speed does not vary significantly from one second to the next. Knowing the vehicle's location, the roaming decision primitive uses the list of road segment locations to determine in which road segment the vehicle is located. It then uses the handover sequence to check if the vehicle has reached a handover location. If it has, the primitive has two options. Either the MS is already connected to the AP corresponding to the current location in the handover sequence. In this case, the roaming decision primitive ends. Alternatively, the MS is not connected to the AP corresponding to the current location in the handover sequence, because it is still connected to the previous AP in the list or it is disconnected. In this case, the primitive runs a pre-connection mechanism consisting of probing the selected AP. If there is no answer after three Probe Requests, the primitive stops probing and waits for 20 ms (a reasonable bound for the Probe Response delay). If the MS receives a Probe Response from the selected AP, the primitive calls the connection primitive, which triggers the roaming process. This mechanism is used because the MS does not perform AP scanning, and thus the MS does not know if the selected AP is reachable when it decides that the handover needs to be triggered. If the selected AP cannot be reached, the roaming process will disconnect from the current AP without attempting a re-connection. The probability of such event seems negligible but as shown by Mahajan et al. [Mahajan 2007], there are numerous obstacles and sources of temporary interference (such as heavy vehicles) in the urban environment that can cause "grey periods", i.e., short periods of very poor RSS. The notion of the grey period should be distinguished from the notion of"white areas" which correpond to areas that are not covered by any AP.

3.3.5 Improving MROAD

MROAD addresses the following two objectives: (i) reducing handover impact, and (ii) improving AP selection. Reducing handover impact is achieved by eliminating the scanning phase known to be the most time-consuming phase of the handover. In addition, AP selection is improved with the help of prior knowledge of the network that allows the MS to trigger the handover as soon as the candidate AP becomes the best AP. However, as there is no scanning, the MS must rely on static information for AP selection. Consequently, it does not take into account changes in the network deployment. For instance, MROAD does not offer any solution if the AP changes its channel or becomes unavailable. In addition, it relies on the unrealistic assumption that the vehicle's route is known in advance in order to pre-compute the handover location.

Nevertheless, MROAD suggests that it is possible to eliminate scanning and that a scan-free handover technique can be efficient in selecting an AP.

3.4 Adding vehicle mobility to the AP selection process

3.4.1 Basic principles

The solution described here overcomes the main drawback of MROAD, as it does not require prior knowledge of the vehicle's route. The COntext-aware Predictive handovER (COPER) technique includes an AP selection module that predicts vehicular mobility in the near future such that the MS is able to use prior knowledge of the network to perform the handover at the most appropriate location. Mobility prediction is provided by a Direction Detection Module (DDM) that analyzes the trajectory of the vehicle. It removes the need for scanning even when the driver does not provide his/ her destination nor the route he/ she plans to take. In order to detect direction and select APs, COPER uses the same information as MROAD with the exception of the vehicle's route: road topology, the location of APs and their modeled RSS are all stored in a Context DataBase (CDB).

Like MROAD, COPER is based on a roaming decision technique that always provides the best connection but avoids scanning. Selection is based on prior knowledge of the road topology, the locations of APs and a model of their RSSs. Using this information, the MS is able to determine its location on the road network and use the modeled RSSs of nearby APs to choose the one that provides the best signal in the new location.

As illustrated in Figure 3.14, COPER is composed of three main modules intended to provide seamless transitions between the best candidate APs regardless of the vehicle's route. To this end, the DDM anticipates the direction of the vehicle using a fuzzy logic-based trajectory analysis that takes into account road characteristics that have been stored in a pre-computed CDB. To predict direction, the AP Selection Module (ASM) uses the simulated AP RSS stored in the CDB in order to identify the best APs and the best locations to trigger handovers. Finally, the Connection Module (CM) detects when the vehicle is entering a handover area and attempts a connection by probing the selected AP and associating with it. The following sections describe the construction of the CDB and the COPER modules.

3.4.2 CDB construction

COPER decisions are based on road topology and the modeled RSSs that are stored in a CDB. To build this CDB, we need to gather multiple, publicly-available sources of information and organize them. In this work we use two notions to characterize road topology. The first is the Road Portion (RP) defined as the portion of road bounded by two intersections or a dead end. An RP is characterized by an azimuth α and the coordinates of the starting point D and the ending point E. The second notion is the Road Segment (RS), which is a sub-section of the RP. Each RP is composed of a set of RSs with equal size (around 5 m). Note that both RP and RS are oriented, meaning

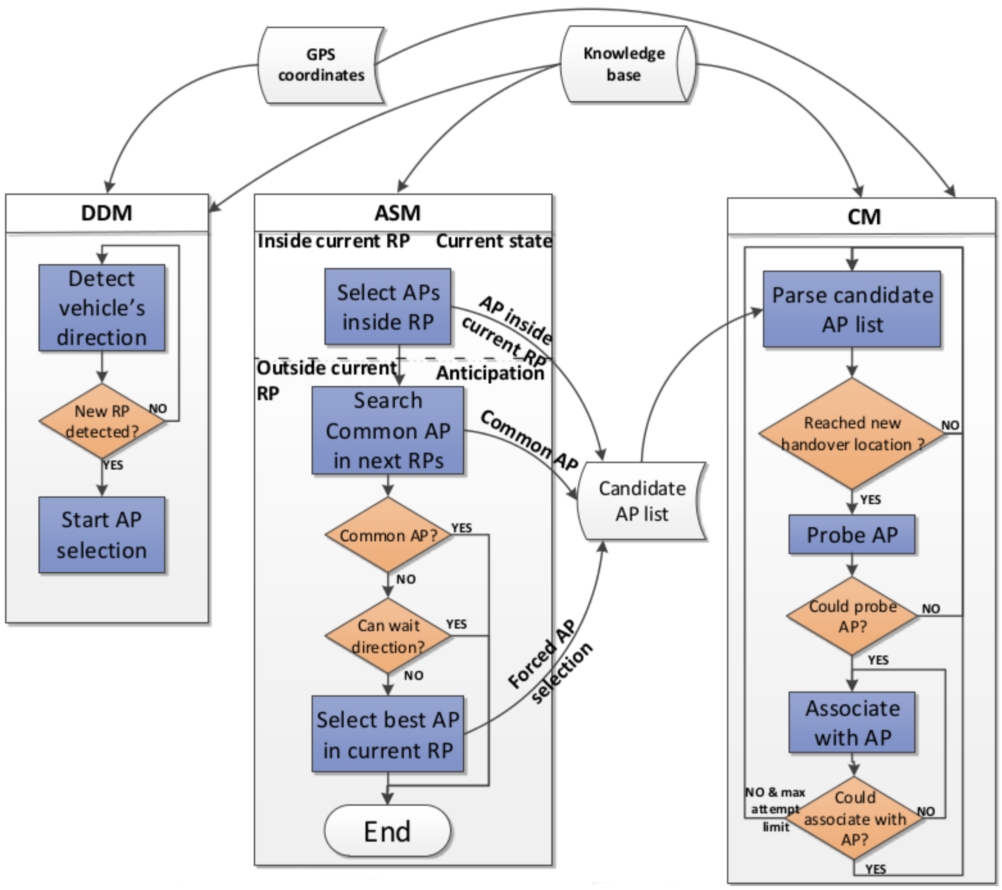

Figure 3.14: Description of the three main COPER modules

that a two-way street is composed of two RPs. The CDB structure consists of five tables. Three describe road topology: the first lists RPs including their boundaries (in terms of GPS coordinates), size and RSs; the second lists the RSs including their boundaries, size and the RP they belong to; and the third lists the relations between RPs. The AP table stores AP locations, SSIDs, frequencies and the RP it belongs to. Finally, the last table contains the RSSs of the listed APs on each RS obtained through simulation.

To build the CDB, COPER uses the same source of information and simulation frameworks as MROAD. Road topology is extracted from the OpenStreetMap[4] database and the list of APs and their locations from an open database provided by the local administration. Note that we also extracted the urban topology for COPER. Next, we used SUMO and Qualnet with the Vergilius framework [Giordano 2010a] to compute RSSs. In addition to CORNER, Qualnet takes into account urban topology in the signal propagation computation, in order to provide more realistic RSS simulations. The database contains the name and the GPS coordinates of each AP in the municipal network. It is used to feed a website that maps the different facilities provided by the municipality[5]. The complete list of SSIDs was provided by the network operator.

A partial visualization of the CDB is shown in Figure 3.15. Based on signal propagation modeling, the CDB can provide information about the number of APs broadcasting with a minimum required RSS (-85 dBm in Figure 3.15).

3.4.3 Direction detection module

3.4.3.1 Description

COPER relies on the characteristics of vehicular mobility as described in Section 3.1.1. In particular, the vehicle's trajectory is guided by the road infrastructure and driving regulations, which makes predictions much easier. Based on this, and the notion of the RP previously defined, COPER considers the following mobility scenario. When a vehicle has been located in a RP, COPER assumes that the vehicle continues on this RP until it reaches the next intersection i.e., the intersection that ends the current RP. COPER then considers the RSS of nearby APs in the current RP and beyond, in the RPs accessible at the next intersection. Direction predictions are performed when the vehicle approaches the next intersection. This attempts to establish as soon as possible which RP the vehicle is entering, based on its trajectory.

The DDM has two roles: (i) it locates the vehicle in an RP, and (ii) attempts to predict the RP the vehicle is about to enter. It relies on GPS location updates and the road topology provided by the CDB. This information is critical in order to trigger the handover at the right location; in particular, when the vehicle turns at

[4] http://openstreetmap.org
[5] http://www.topographie.lu

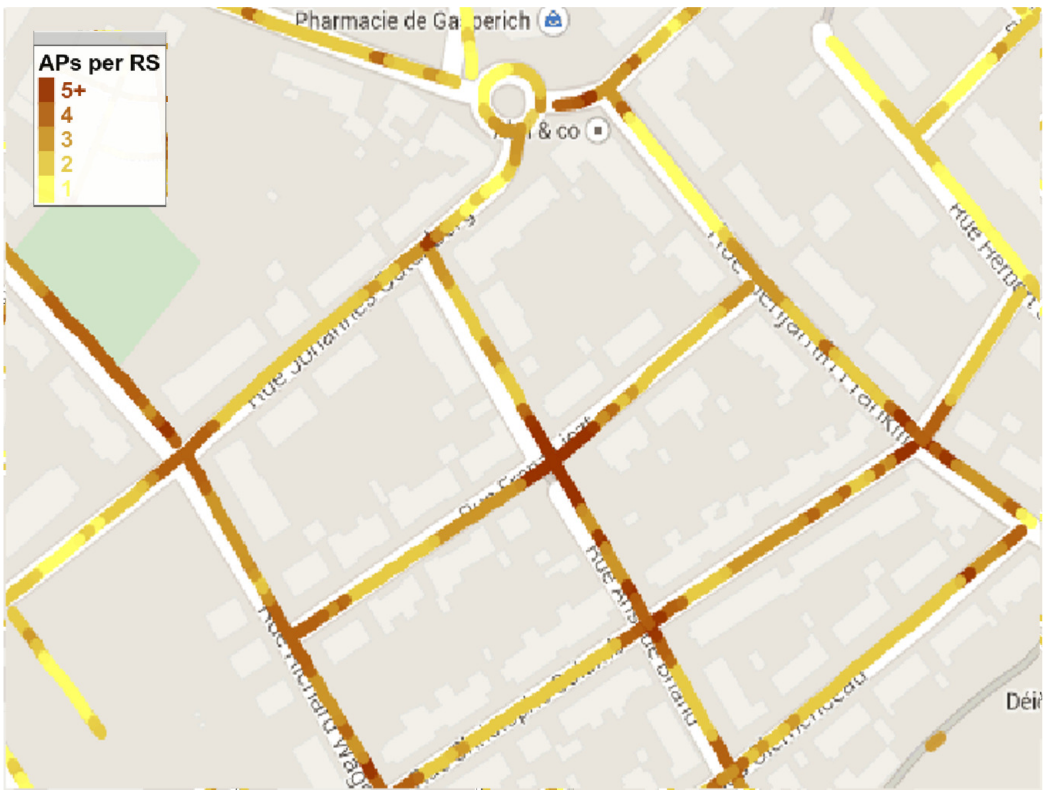

Figure 3.15: AP density heatmap based on the CDB

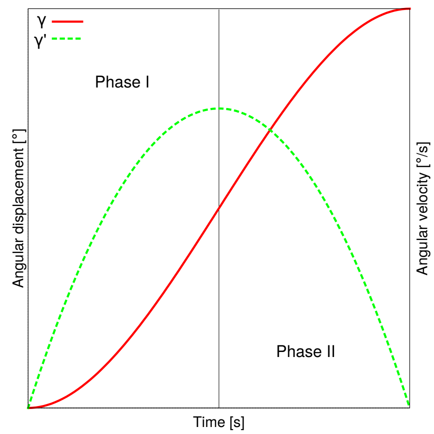

Figure 3.16: Model of angular displacement during a turn.

an intersection and is about to leave the LoS of the current AP. This situation often leads to a significant decrease of RSS from the current AP, causing packet loss or even complete disconnection. The DDM distinguishes between a vehicle that continues in the same direction from one that turns at an intersection. The first case corresponds to only one candidate RP. Therefore, we consider that the vehicle continues straight on from the current RP to a candidate RP, if the azimuth difference between RPs is less than 10 °. Otherwise, the candidate RP implies that a turn has been made. This threshold issues from the observation that intersections always comprise only one adjacent RP with an azimuth difference lower than 10 ° which is not the case if one considers higher thresholds. The next direction is found by comparing the current trajectory with the characteristics of the candidate RPs (i.e., the RPs starting at the next intersection). The current trajectory is defined by the two most recent vehicle locations l_t *and* l_{t-1}, and its bearing β. The next direction detection starts when the vehicle enters the direction detection area (15 m before the end of the current RP). During this phase, the DDM computes the angular shift γ such that $\gamma = |\beta - \alpha|$. The DDM concludes the vehicle has entered a new RP when γ falls below 20 ° and the distance between the vehicle and the RP entry point decreases (Equation 3.5). This threshold has been empirically obtained during a preliminary study. It results from a trade-off between the risk of a wrong detection and the objective of early detection.

$$\begin{cases} \gamma < 20^\circ & (a) \\ d(l_t, D) < d(l_{t-1}, D) & (b) \end{cases} \tag{3.5}$$

If the DDM has not detected any turning 10 m after the end of the current RP, it concludes that the vehicle has continued straight ahead. On the one hand, this rule negatively impacts the performance of the DDM in terms of detection anticipation but, on the other hand, it allows the DDM to detect a greater number of turning events. In addition, the early detection that the vehicle has not turned is considered less critical than the detection of a turning event, due to the fact that the MS is more likely to stay longer in the Line of Sight (LoS) of the current AP when it does not turn.

Note that as the vehicle turns, its bearing starts to coincide with the azimuth of the next RP, and it is thus too late to make any predictions. Therefore, in order to better anticipate the turn, we augment the DDM with an Fuzzy-based Turning Detector (FTD). The FTD uses a fuzzy logic system to evaluate the probability that the vehicle is initiating a turn. The FTD relies on a simple model that characterizes a turn based on the angular displacement γ. In our model, γ is characterized by a two-phase change (phase I: initialization; phase II: stabilization) as depicted in Figure 3.16. Initially, both γ and the angular velocity $\gamma\prime$ are zero, i.e. the vehicle is moving in the same direction as the current RP. In phase I, γ increases from zero until it reaches a maximum, when the turning phase is initiated. In phase II, the angular velocity decreases until there is no more angular displacement. Entering phase II initiates trajectory stabilization.

The input variables of the fuzzy system are: vehicle speed, its acceleration (computed as the speed first derivative) γ, and its angular velocity $\gamma\prime$ as the derivative of γ. In all cases, we consider trapezoidal membership functions. The output variable, the *turning metric*, is a value in the interval $[-180; 180]$ interpreted as the most likely angle of the vehicle's trajectory. We define seven linguistic variables for the turning metric: one for the non-turning case, and three (low, medium and high) for each turning direction (i.e., negative values imply a left turn, while positive values imply a right turn). The FTD considers a set of rules that can be summarized as follows: (i) there is no turn (i.e., the turning metric equals zero) if γ is low or vehicle speed is high; (ii) if vehicle speed is low and there is negative acceleration, the turning likelihood is high; (iii) if the angular rate is non-zero and γ is increasing, the turning likelihood is high (phase I in Figure 3.16); (iv) if the angular rate is null or γ is decreasing, the turning likelihood decreases (phase II in Figure 3.16); and finally, (v) if the vehicle accelerates quickly from a low speed, the turning likelihood increases. Note also that the angular rate is used here as a filter, as it indicates whether or not the turning manoeuvre continues (see fourth and fifth rules). In order to select the candidate RPs we evaluate the relative angle between the azimuth of the current AP and each candidate. We define seven groups: one for the RP that implies the vehicle is continuing straight on (relative angle below $10\,^{\circ}$), and three groups (for turning right and left) for RPs that are linked to a low, medium or high turn for relative angles between $10\,^{\circ}$ and $45\,^{\circ}$, $45\,^{\circ}$ and $110\,^{\circ}$ and higher than $110\,^{\circ}$ respectively. Note that the case of two candidate RPs that are classified in the same group is infrequent and considered out of the scope of this study. Unlike the previous approach, the FTD can anticipate the vehicle's direction. However, it can only be used to detect turning; thus, the DDM combines these two approaches and identifies the one that provides earliest detection.

3.4.3.2 How far can the DMM predict?

	Median detection	Detection $< 10\,\mathrm{m}$
DMM	12 m	45%
DMM + FTD	-7 m	92%

Table 3.2: Turn detection performance of DDM and DDM with FTD

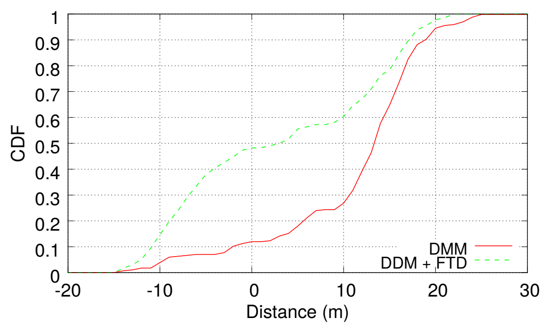

Figure 3.17: CDF of anticipation distance

Here we investigate DDM performance. The route for this experiment is 6.5 km long and comprises 62 intersections where the vehicle performs 43 turns. This scenario includes multiple types of intersections, implying different mobility patterns (i.e., different speeds and trajectories). We use a 4 Hz differential GPS receiver in order to have frequent updates of the vehicle's trajectory and decrease detection latency. The experiment was repeated multiple times in order to obtain a representative overview of DDM detection performance. As mentioned earlier, the most critical part of direction detection is turn detection: when the vehicle turns at an intersection, the MS is more likely to leave the LoS of the current AP, and consequently there is likely to be a significant decrease in current AP RSS. Table 3.2 highlights performance in terms of turning detection for the standalone DDM and the DDM combined with the FTD. We use the distance between the detection location and the end of the current RP as the performance metric. This distance is negative when the detection occurs before the vehicle reaches the end of the current RP (anticipated detection) and positive otherwise. We also evaluated the false negative rate, which corresponds to the case where the DDM does not detect that the vehicle is turning or selects an incorrect RP, and false positives that correspond to the case where the system detects a turn but the vehicle is not turning. From 11 experiments, we obtained seven false positives and seven false negatives compared to 129 successful detections (95%) for both standalone DDM and DDM augmented with FTD. This suggests that the addition of the FTD does not have an impact on the detection success of the DDM. However, FTD enhances median detection anticipation by 19 m. In addition, 92% of detections occur within 10 m versus 45% for the standalone DDM. If we focus on overall direction detection (straight RPs and turns), as depicted in Figure 3.17, we observe a median detection anticipation of 10 m when DDM is augmented with FTD (3 m vs. 13.2 m).

The Direction Detection Module (DMM) predicted vehicle direction at half of the intersections and early detection otherwise. This ability to anticipate allows the ASM to adapt AP selection to the vehicle's route in real time, in particular when the vehicle is about to turn. In turn, the ASM can detect when the RSS of the current AP may decrease in the short term and anticipate the selection of the next AP such that the handover is triggered before the current AP's RSS falls dramatically.

3.4.4 COPER AP selection

The ASM tries to find the AP that provides the best network connectivity based on the vehicle's direction. As mentioned earlier, the MS does not know the complete route in advance. This mean that AP selection must be triggered every time it enters a new RP. The MS first evaluates possible handovers to candidate APs located in the current RP. It then evaluates handovers to APs located in the next RPs and makes an anticipated handover decision if possible. The selected APs are then stored in the *candidate AP list* used by the connection module (cf. Figure 3.14).

Handover evaluation and localization The ASM uses the same handover evaluation methods as MROAD. For each candidate AP it builds the handover quality metric that is used in AP selection. Figure 3.18(a) shows the handover signal quality Q_S as the maximum of the curve delimiting the pattern-filled area and the handover location H_S (the RS corresponding to HS_S). In addition, it computes the handover location as the location where the RSSs of both the current and the candidate AP are simultaneously maximized. A detailed description of the handover location computation can be found in Section 3.3.3.

AP selection process Figure 3.18(a) shows a sample of simulated and real RSSs of two APs called respectively AP_i (current AP) and AP_j (candidate AP). Note that the CDB only provides the modeled RSS to the ASM, while the sensed RSS is only used to evaluate ASM performance.

AP selection is triggered every time a new RP is detected by the DDM. The ASM

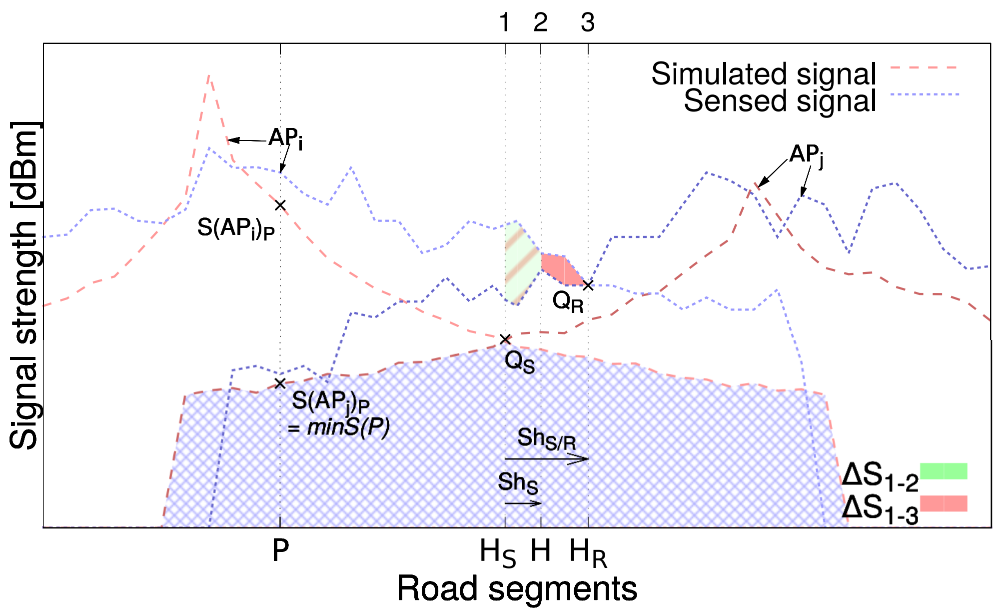

(a) Case 1

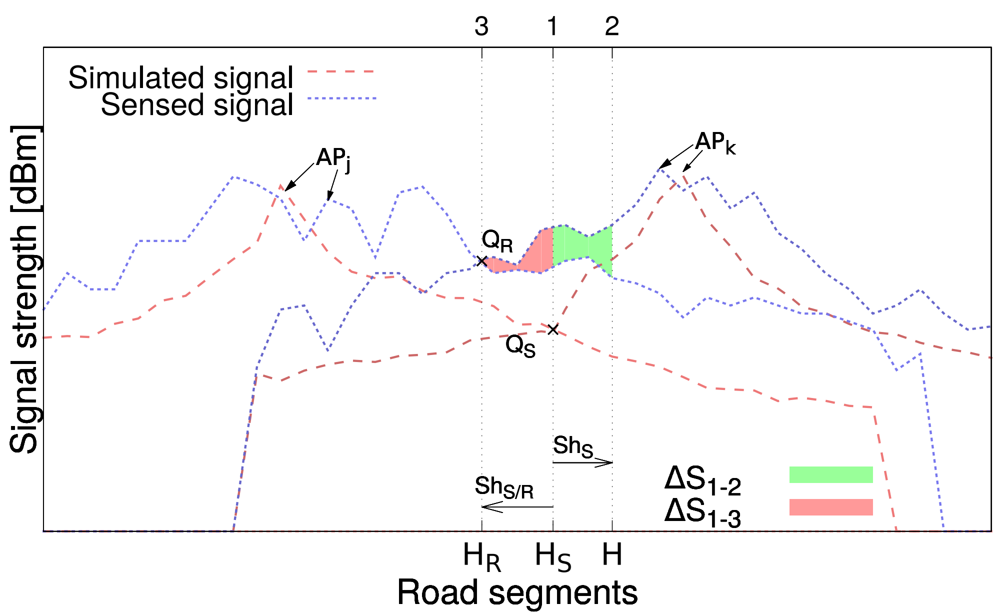

(b) Case 2

Figure 3.18: Sample of simulated and sensed RSS

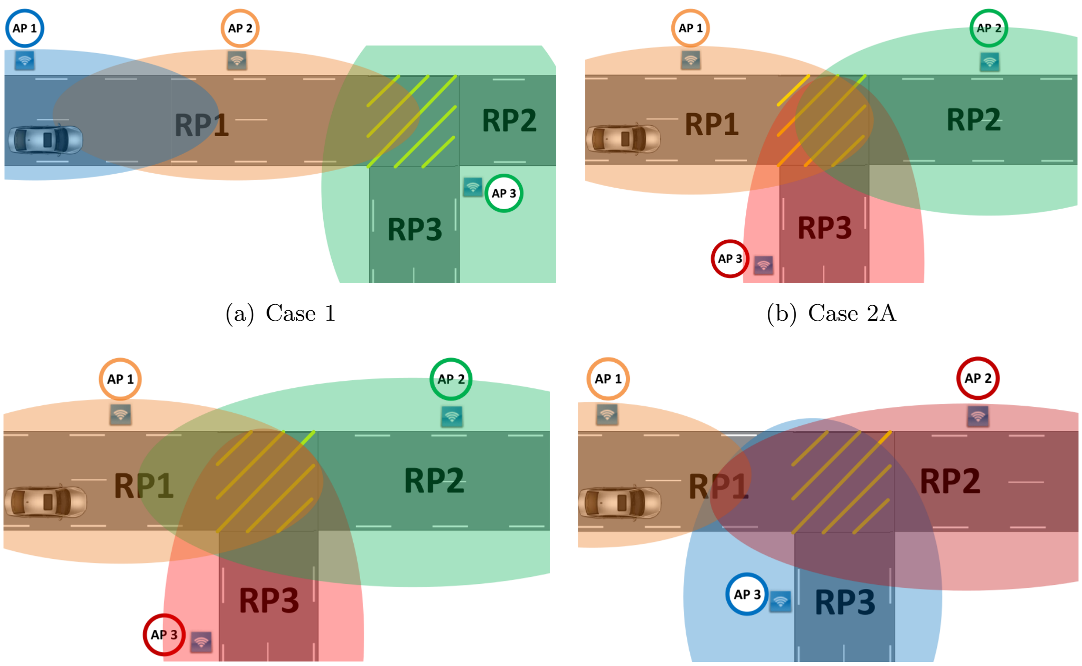

(a) Case 1 (b) Case 2A

(c) Case 2B (d) Case 3

Figure 3.19: Different cases handled by the ASM

automatically selects the APs in this new RP and explores potential handovers to an AP located in the next potential RPs. The ASM classifies the network beyond the selected RP into three groups that cover all possible scenarios. These scenarios are depicted in Figure 3.19. First, the ASM looks for a single AP that is shared with the next RPs (Figure 3.19(a)). If such an AP is found, it is selected, the handover location is computed and AP selection is resumed the next time the DDM detects a new RP. If not, the ASM analyzes the RSS of the current AP. If this RSS remains sufficiently high to the end of the RP, AP selection ends (Figure 3.19(b) and Figure 3.19(c)). If not, the ASM forces the selection of another AP located in the next RPs (e.g., AP2 or AP3 in Figure 3.19(d)). Note that when an AP is selected, the ASM computes the handover location (H_S) as described before and provides the CM with the 3-tuple $\{AP_{from}, AP_{to}, H_S\}$ where AP_{from} is the currently selected AP and AP_{to} is the next selected AP.

A comparison of the AP selection processes of COPER and MROAD shows that removing the assumption that the MS knows the vehicle's route significantly increases complexity. However, it does not mean that AP selection in COPER is suboptimal compared to MROAD.

We can identify two situations where the MS is not associated with the best AP. The first is due to an incorrect estimate of RSS, leading to situations where a handover location that is based on the simulated RSS (H_S) does not correspond to the optimal handover location in reality (H_R). The spatial shift between those two handover locations is represented by $Sh_{S/R}$ (see Figures 3.18(a) and 3.18(b)). $Sh_{S/R}$ is considered positive if H_R is located after H_S and negative otherwise. For a given period, the RSS of the current AP is lower than the RSS of the best AP. This RSS difference is called ΔS_{1-3}. The second suboptimal AP selection case is due to short-term predictions of vehicle's direction. As the DDM often detects the next RP at the end of current RP, if the current AP provides a good signal, the ASM delays AP selection until the end of the current RP. This sometimes leads to suboptimal AP selection. In case 3 shown in Figure 3.19(c), AP2 and AP3 provide higher RSS at the end of the current RP. In this case, AP selection is suboptimal as the ASM does not select the best available AP. This situation is shown in Figures 3.18(a) and 3.18(b)), where the actual handover location H does not correspond to best option H_S. The spatial shift between the two is denoted Sh_S and the RSS difference as ΔS_{1-2}. Sh_S is always positive as H can never be ahead of H_S.

3.4.5 Connection module

The CM triggers a new association with the selected AP at the handover location given by the ASM. This module only takes as input the selected AP list and the current vehicle location without performing any AP scanning. Figure 3.14 describes the connection procedure and the CM. First, the CM parses the selected AP list by comparing the selected AP handover location with the current location of the vehicle. If the handover location of a selected AP is reached or has been passed, the CM starts the procedure for connection to this AP. The connection procedure consists of two phases. First, since the CM does not perform any scanning, it does not know if the selected AP is actually available or not (as the fact that the selected AP is listed in the CDB does not guarantee it is the case). As a result, the CM sends it a Probe Request. If the CM does not receive a Probe Response after three attempts, the selected AP is considered unavailable. The ASM then begins AP selection again, ignoring this unavailable AP. On the other hand, if a Probe Response is received from the selected AP, the CM initiates the association process. This process is the standard IEEE 802.11 authentication and association, i.e., the exchange of four messages (authentication request/ response and association request/ response) with the selected AP. If this phase fails, the CM retries a connection to the selected AP.

3.5 Conclusion

Based on the study of IEEE 802.11 communications under vehicular mobility, we identified two main improvements to vehicular communications. The first reduces the length of the handover (where scanning latency is the main problem); the second enhances AP selection which can cause sudden link disruption.

As a first step, we studied an optimized scanning-based approach that consisted in using a second interface base on the make-before-break paradigm. Although this scanning-based approach successfully addressed the first issue, it did not provide enough information to improve AP selection.

Next, we proposed MROAD, a scan-free approach based on prior knowledge of network topology and the location and RSS of APs. By removing the scanning phase, this approach significantly decreased the length of the handover based on prior knowledge of the network and enables the MS to select the best APs on its route.

Finally, we proposed COPER that extends MROAD by introducing vehicular mobility in the AP selection and avoids the need to have prior knowledge of the vehicle's route.

Chapter 4

Evaluation

Contents

4.1 Introduction

This chapter evaluates the practical application of the three approaches described in the previous chapter. We begin with a presentation of the experimental testbed, a residential neighborhood in the city of Luxembourg where HotCity, an IEEE 802.11 hotspot network is deployed. This studies performed in this real-life testbed make it possible to compare the new solutions against legacy, scanning-based approaches and to identify their limitations. Finally, we evaluate the extent to which the three proposed solutions meet requirements with regard to the two objectives defined in this thesis: (i) eliminating handover latency and (ii) improving AP selection by taking into account vehicular mobility.

4.2 Testbed description

In order to evaluate the performance of the solutions presented in this thesis, we performed a set of experiments using the HotCity network, a metropolitan IEEE 802.11 network deployed over a large part of the city of Luxembourg. HotCity is composed of around 500 Cisco Aironet APs[1] for outdoor use, each embedding an IEEE 802.11b/g/n radio for client wireless access and an IEEE 802.11a radio for mesh interconnection between APs. In order to guarantee fair usage of the available bandwidth, the per-client downlink data rate was limited to 2 Mbps until 2014. Since then the limit has been raised to 25 Mbps. The testbed was the Gasperich neighborhood, a residential area of Luxembourg city. We choose this area, rather than the city center in order to have a more fluid vehicular mobility pattern, without being affected by traffic jams. In this area, the speed limit is 30 km/h and vehicles have to give way to other vehicles coming from the right. Hence, during the experiments the car often stopped. Although traffic varied during the day it remained fluid. A map of the hotcity AP deployment is available online[2].

In order to characterize the testbed, we used Wi2Me [Castignani 2012b], an Android wireless network analyzer for IEEE 802.11 and 3G networks. Wi2Me gathers the RSS of nearby APs, their channel and their capabilities in terms of security. We show in Figure 4.1(a) the Cumulative Distribution Function (CDF) of the number of APs per scan both for all the discovered networks and for HotCity APs. The distribution of the discovered APs reveals a median of 9 APs per scan, denoting very dense deployment. This high density is mainly caused by residential IEEE 802.11 APs belonging to personal networks. It is important to note that for all scan results, at least one HotCity AP was discovered, which means that HotCity completely covers the path considered in this study. This is the case in all the experiments presented in this section. Regarding AP quality, the CDF of the RSS in Figure 4.1(b) shows that HotCity APs have a higher RSS than other APs (e.g., they provide a median of −82 dBm against −87 dBm for all APs). This difference in RSS is due to the fact that personal IEEE 802.11 networks are indoor whereas HotCity has been designed for outdoor wireless access. Note that, as shown in Figure 4.2(a), an RSS lower than −80 dBm significantly degrades the QoS. As a result, mobile users may encounter poor connectivity. Note also that in roughly 10 % of scans, only one HotCity AP is discovered. This may lead to low throughput if the MS only considers the HotCity network.

Finally, the channel distribution in Figure 4.2(b) shows that APs are mainly distributed over channels 1,6, and 11, which are three non-overlapping 2.4 GHz channels. This is a typical configuration in current urban deployments [Eriksson 2008, Giordano 2010b, Castignani 2011b].

[1] http://www.cisco.com/c/en/us/products/wireless/access-points/index.html
[2] https://www.hotcity.lu/www_laptop_en/about/wifi_coverage

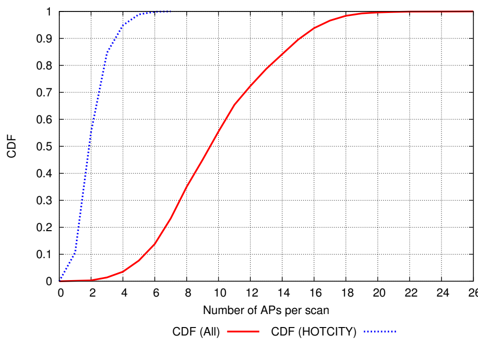

(a) Number of APs per scan

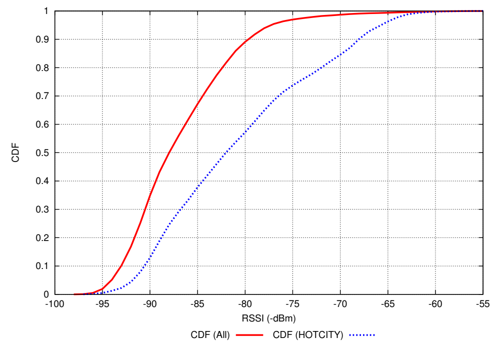

(b) CDF of RSS

Figure 4.1: HotCity deployment characteristics

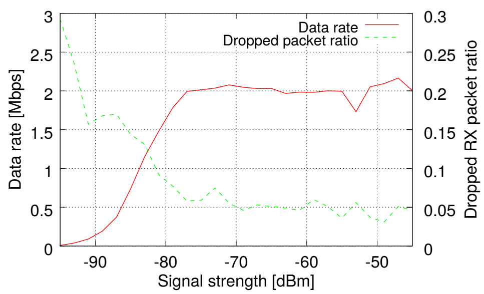

(a) Relationship between RSS, data rate and dropped packets

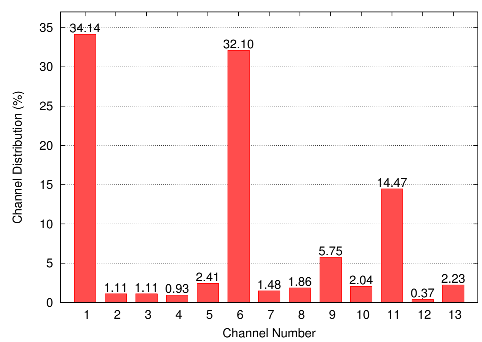

(b) Channel distribution

Figure 4.2: HotCity deployment characteristics (cont.)

The experimental route was chosen such that the mobile user always receives a sufficiently high RSS. For instance, Figure 4.3 highlights the number of APs and the maximum RSS detected along the experimental path used in COPER. The distribution of the number of APs shows that there is no disconnected area. Also, in all cases the highest RSS is always greater than -80 dBm. Consequently, if the MS efficiently selects APs, it will always be associated with a HotCity AP and have Internet connectivity. Note that a detailed characterization of the HotCity network is available in [Castignani 2012a].

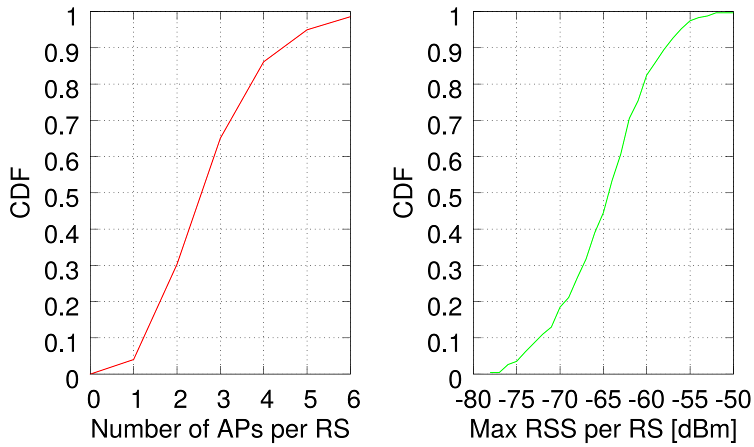

Figure 4.3: CDF of the number of APs per RS and mean data rate per RSS

4.3 Using a second interface: the make-before-break approach

4.3.1 Experimental setup

In order to analyze the performance of the static and dynamic algorithms we performed a set of experiments to quantify the performance of a second interface (used to optimize the handover) compared to a legacy, single-interface handover implementation. To this end, we configured a laptop embedding two IEEE 802.11 USB interfaces (Edimax EW-7711USn) connected to two 3 dBi gain external antennas placed on the roof of a Renault Twizy electric car. We placed the two antennas more than 80 cm apart in order to avoid interference that occurs when two IEEE 802.11 interfaces that are located close to each other are used simultaneously. This issue has been studied in [Nachtigall 2008, Robinson 2005]. The laptop implemented the Wireless Extensions (WEXT) interface for Linux over the rt2680 driver in the kernel space. For the dynamic algorithm, we modified the laptop configuration in order to assign the same MAC address to both interfaces.

The path taken for the experiments was 1.8 km long and included several intersections, one roundabout, and one traffic signal.

The experimental scenario is as follows. At the point of departure, while the car is still stationary, a 500 kpbs User Datagram Protocol (UDP) flow download is initialized from a remote server located on the campus of the University of Luxembourg. The experiment starts when the car starts moving. During the experiment, we recorded packets received using *tcpdump*, the output of the WPA-supplicant debug, and the GPS location of the car. Note that we use DHCP to obtain an initial IP configuration before the experiment began and retained it for the duration of the experiment. When using the dynamic algorithm, we recorded received packets on both interfaces. The two traces were merged and results were analyzed offline. We used *editcap*, part of the wireshark suite[3], to remove duplicate packets observed when using the dynamic algorithm. Duplicate packets occur because both interfaces have the same MAC address. It is important to remove them during the evaluation in order to measure the actual packet delivery ratio.

4.3.2 Evaluation of throughput and connected time

In this section we analyze the connectivity performance of the MS. We consider three different handover approaches: a Single Interface (SI) approach, and the static and dynamic algorithms implementing a second interface. First, we focus on Layer-2 (L2) connectivity duration, i.e., how long the MS is associated with an AP. For the SI approach, we consider the connected time as the time that elapses between MS association and when a scan is triggered due to low RSS. When the MS uses two interfaces and implements the static algorithm, the connected time is the time that elapses between the instant the MS is successfully associated, and the instant it triggers a new association to the candidate AP. In the case of the dynamic algorithm, the MS is considered to have L2 connectivity if at least one of its interfaces is associated to an AP. Figure 4.4(a) shows the minimum, average and maximum percentages of the time the MS remains connected at L2. With the SI approach, the L2 connection lasted between 52.3 % and 69.7 % of the time, with an average of 62.5 %. Compared to the SI approach, we find that the additional interface offers a significant improvement in terms of L2 connectivity. The static algorithm provides L2 connectivity between 71.3 % and 84.7 % of the time, with an average of 79.2 %, which is 17 % higher than the SI algorithm on average. Finally, the dynamic algorithm outperforms both SI and the static algorithm. It is connected at L2 for an average of 97.9 % of the time (i.e., between 95 % and 99.3 %). The dynamic algorithm achieves almost continuous connectivity in all experiments. The second metric we propose is the Packet Delivery Ratio (PDR), (i.e., the number of packets successfully received by the MS) for a 500 kpbs UDP downlink flow. This is shown in Figure 4.4(b), and allows us to estimate the quality of Layer-3 (L3) connectivity while moving. The downlink flow used was not particularly intense, but rather illustrates loads produced by a typical multimedia connection. Figure 4.4(b) shows the minimum, average and maximum PDR for each handover algorithm. The same trends are found as for the L2 connection: the SI approach has the lowest PDR (68.89 % on average), while the two-interface approach outperforms it with 79.9 % for the static algorithm, and 95 % for the dynamic algorithm.

[3]https://www.wireshark.org/docs/man-pages/

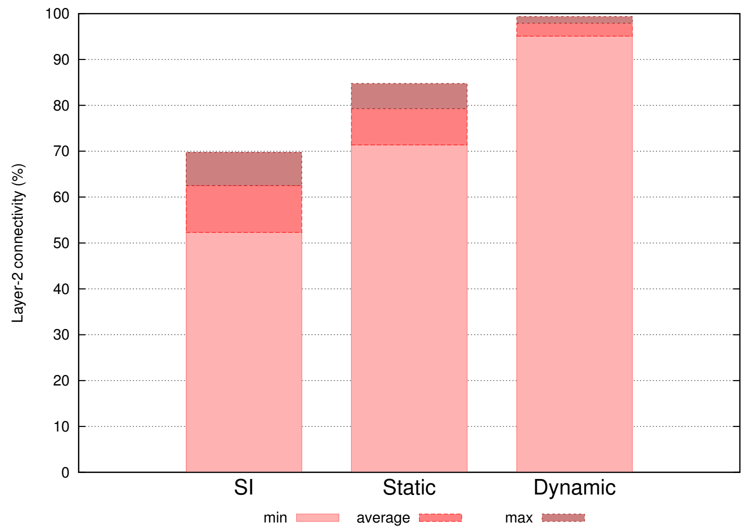

(a) Layer-2 connectivity

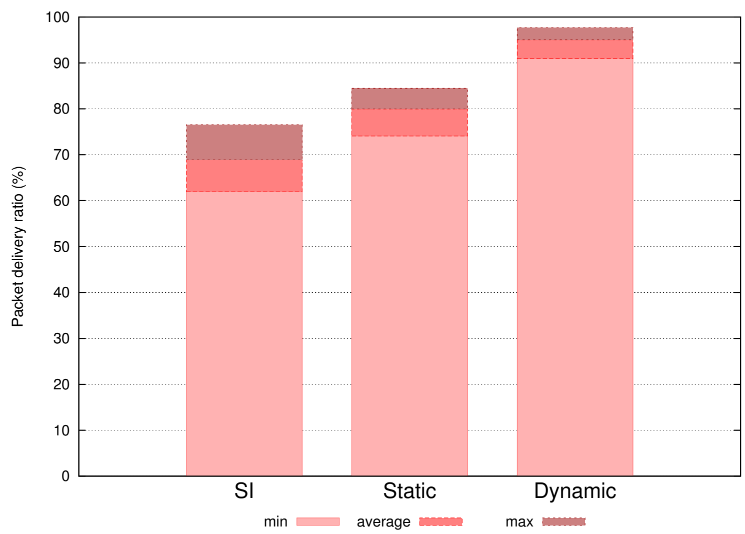

(b) Packet delivery ratio

Figure 4.4: General results

4.3.3 Impact of handovers on performance

In order to investigate these results in more details, we took a representative experiment of each handover approach (SI, static and dynamic algorithms) in order to compare their performance and impact. Note that, as the speed of our vehicle is not strictly constant, the duration of the experiments varies between 250 s and 285 s. We chose three

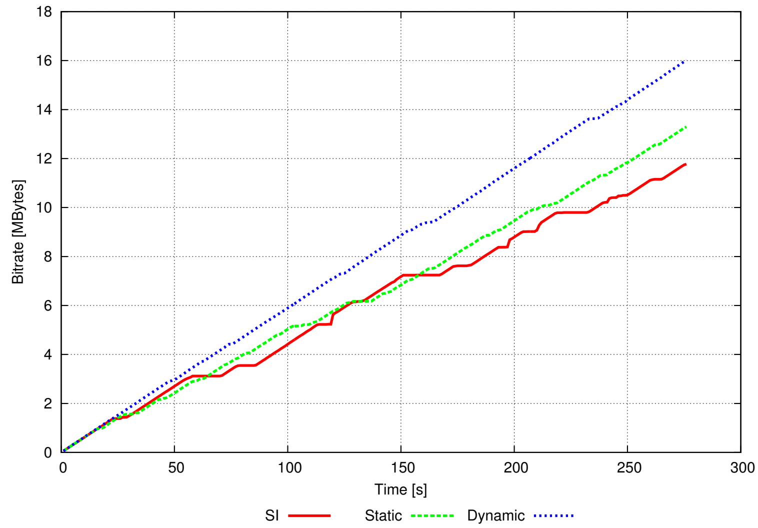

(a) Cumulative downloaded data

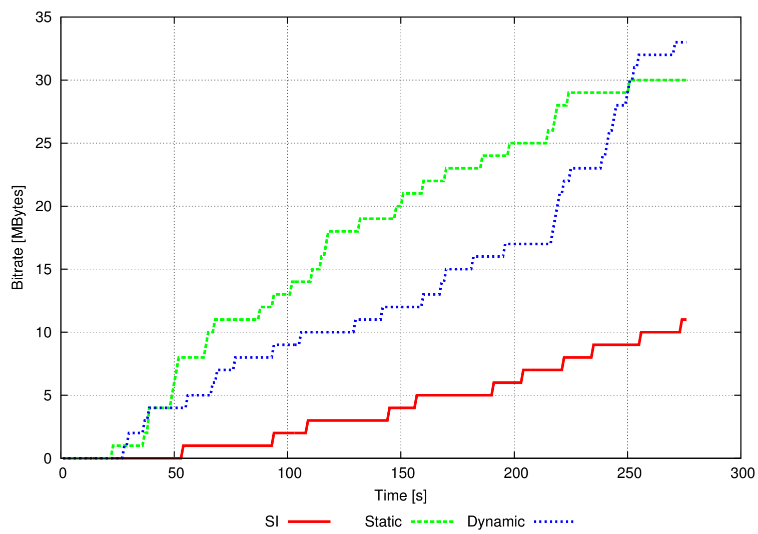

(b) Cumulative number of handovers

Figure 4.5: A representative single run

experiments lasting approximately the same amount of time. Figure 4.5(a) shows the changes in the number of bytes downloaded for the three approaches for the duration of the experiment. Note that for all the approaches, throughput is identical for the first 25 seconds. The first handover occurs at time 25 s. The change in the number of downloaded bytes observed for SI and the static algorithm differ from observations for the dynamic algorithm. This is because the SI and the static algorithms lead to a disconnection during the handover. However, over the long term, the static algorithm downloaded more bytes than SI. When we look at these two curves in detail, it is clear that throughput in the SI approach suddenly falls to zero several times; the same phenomenon is seen for the static algorithm, but lasts for less time. In contrast, throughput growth for the dynamic algorithm is relatively uniform, although there are short interruptions at times 50 s, 160 s, and 230 s.

Figure 4.5(b) illustrates the cumulative number of handovers for each of the three algorithms, and is used to analyze the impact of handovers on connectivity performance. It makes it possible to measure the reactivity of the different algorithms with regard to changes in topology and shows that periods of low SI throughput (shown in Figure 4.5(a)) match the triggering of handovers. We can conclude that the handover impacts the SI approach much more than the static and dynamic algorithms, although the number of handovers is substantially higher for the static and dynamic algorithms. Two reasons can be put forward to explain this: First, note that the implementation does not take into account the “ping-pong” effect, which produces several handovers between the current and the candidate AP when their signal strength is similar. Figure 4.5(b) highlights this ping-pong effect as the number of handovers abnormally increases at certain moments (50s and 120s for the static algorithm, and 220s for both static and dynamic algorithms). A detailed study of these short periods reveals that they correspond to specific locations where there are two HotCity APs providing similar RSS, which is below -75 dBm. The small connection interruptions experienced with the dynamic algorithm are explained by handovers that are triggered close in time due to this ping-pong effect. Liao et al. [Liao 2006] suggest that the hysteresis process can be used to avoid this. Second, since our approach uses a second radio that is continuously scanning, nearby AP information is more frequently updated. This allows the algorithms to rapidly react to changes in topology, i.e., candidate APs are discovered more frequently. The fact that there are more handovers using the static algorithm compared to the SI approach partially explains the fact that the L2 connectivity (see Figure 4.4(a)) of the static algorithm is not as high as expected. However, even this simple solution outperformed the SI approach.

4.4 Replacing scanning by prior knowledge of the network

4.4.1 MROAD

4.4.1.1 Experimental setup

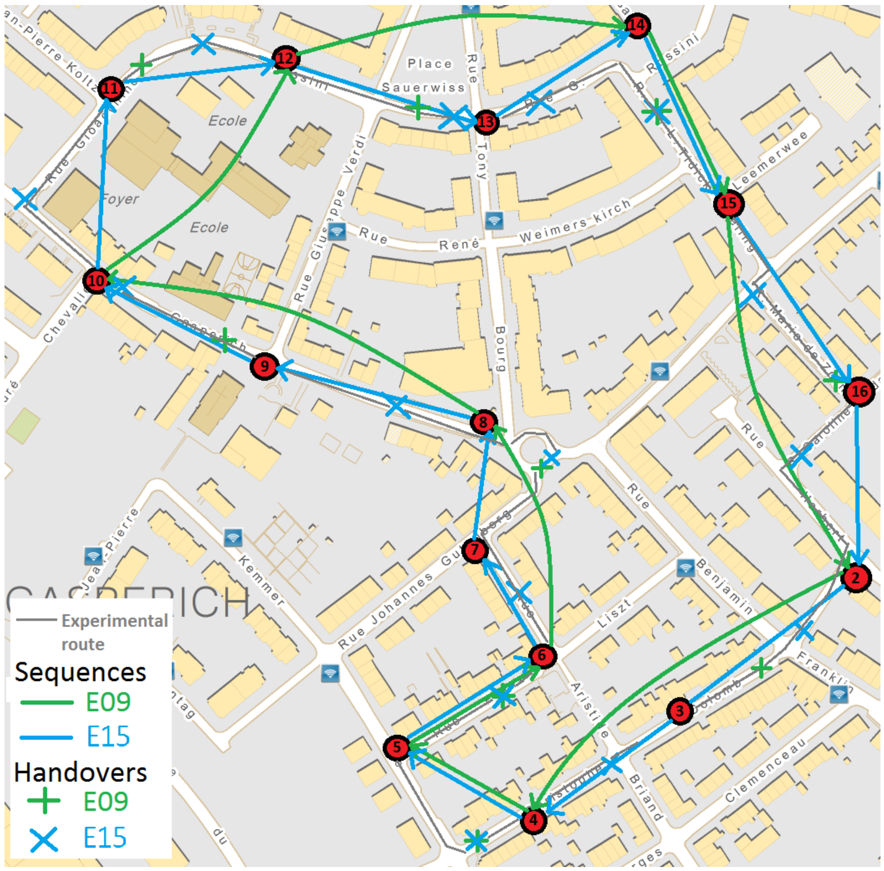

Figure 4.6: Experimental route with shortest and longest sequences

In this section, we evaluate MROAD. MROAD, presented in Chapter 3, Section 3.3, proposes a novel AP selection approach that relies on prior knowledge of the RSS of IEEE 802.11 APs all along the vehicle's route (obtained through simulation) in order to predict optimal handover locations. The experimental route comprises several APs with overlapping coverages. Therefore, the MS has several possible sequences of transitions. In order to identify the sequence with the highest PDR, we consider the trade-off between the impact of increasing the RSS and increasing the length of disconnections (due to more handovers). We formulate this issue as a bi-objective optimization problem (described in Chapter 3, Section 3.3.3).

Label	Sequences on the Pareto-optimal front
E09	2 4 5 6 8 10 12 14 15
E10	2 4 5 6 7 9 11 12 14 15
E11	2 4 5 6 7 8 9 11 12 14 15
E12	2 4 5 6 7 8 9 11 12 14 15 16
E13	2 4 5 6 7 8 9 11 12 13 14 15 16
E14	2 3 4 5 6 7 8 9 11 12 13 14 15 16
E15	2 3 4 5 6 7 8 9 10 11 12 13 14 15 16

Table 4.1: Experimental handover sequences

The study is composed of two parts. In the first, we perform a comparative study of each sequence on the Pareto-optimal front. We then compare our solution with the legacy approach that consists in using a WPA-supplicant with different RSS thresholds (RSS_T:-70 dBm, -75 dBm and -80 dBm). Our earlier work (Section 3.1.1) showed that those three values provided optimal connectivity. The goal of this second study is to evaluate the impact of scanning-free association on connectivity. To this end, we configured a laptop embedding one IEEE 802.11 USB interface (Edimax EW-7711USn) connected to a 5 dBi gain external antenna placed on the roof of a Renault Twizy electric car. The laptop used the *rt2680* driver. Note that for this study, the level of sensitivity *minSens*, defined in Section 3.3.3, was set to -90 dBm. The experiments were performed as follows. When the vehicle was at the departure point, the MS initialized a downlink UDP flow from a remote server at the maximum possible data rate. Note that the MS runs DHCP only once to obtain a unique IP configuration, which is kept for the duration of the experiment (L2 roaming). The experiment began when the car started to move. We recorded received packets using *tcpdump*. Furthermore, we logged the output of the WPA-supplicant debug, the Linux kernel log, the GPS location of the car, together with the signal strength of the connected AP, and finally the quantity of data received and transmitted from *debugfs*. Figure 4.6 shows the experimental route with shortest and longest sequences (this route is identical to that used to study the make-before-break approach). Red circles indicate the location of APs. Note that the AP labeling shown in Figure 4.6 is the same as in the simulation framework. It starts with AP ID 2 because the number ID 1 is reserved for the MS. The experimental run was repeated six times for each parameter in order to obtain representative average values.

4.4.1.2 Results

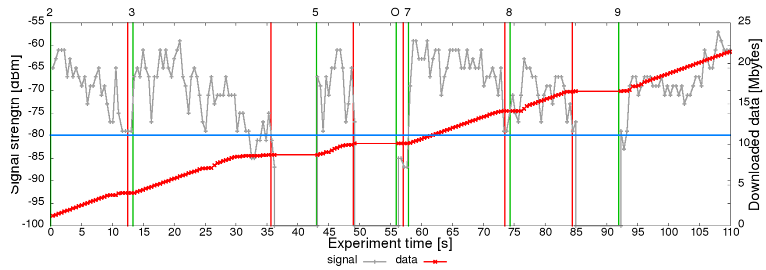

(a) Legacy solution (RSS_T -75 dBm)

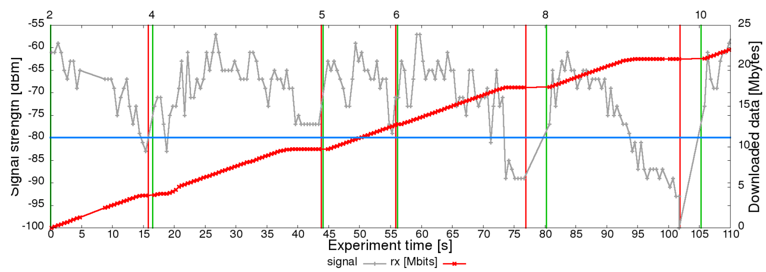

(b) MROAD (sequence length 9)

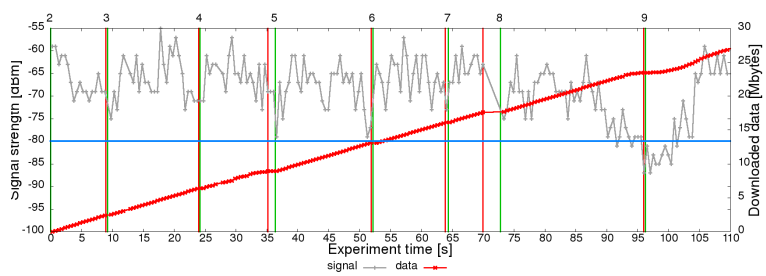

(c) MROAD (sequence length 15)

Figure 4.7: Changes in signal strength and download rate

In this section we analyze several connectivity metrics. As described above, we consider ten cases: the seven handover sequences that appear on the Pareto-optimal front of the two-objective optimization shown in Table 4.1 and the three legacy solutions. For MROAD, handover sequences have different lengths (from 9 to 15 APs) and are called E09 to E15 respectively. For the legacy solutions, the WPA-supplicant was set with three different RSS_T: -70 dBm, -75 dBm, and -80 dBm. These cases are called L-70, L-75, and L-80 respectively.

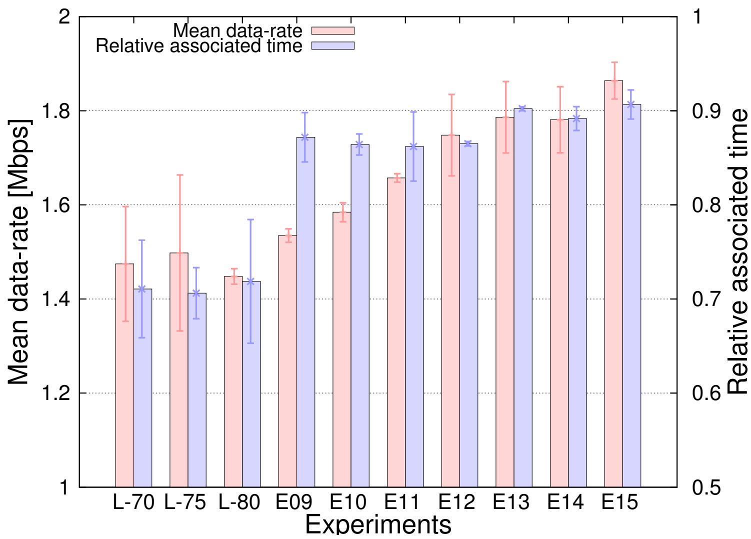

Figure 4.8: Download data rate vs. associated time

Network performance comparison Figures 4.7(a), 4.7(b) and 4.7(c) show representative results of changes in RSS and download rate for the legacy solutions with an RSS_T set to -75 dBm (L-75), and MROAD with the shortest and longest sequences (E09 and E15). The vertical red lines represent disconnections, and vertical green lines show reassociation with a candidate AP whose identifier is indicated above the line. Note that when the AP does not belong to the set of APs in the overlapping AP graph, it is noted as O. We observe that the amount of data downloaded grows faster and more steadily for MROAD compared to L-75, which experienced three plateaus (i.e., zero-throughput periods lasting longer than $5\,s$). After $110\,s$, MROAD was able to download roughly $27\,MB$ compared to only 22 MB with the legacy solution. Despite an additional handover (7 vs. 6), the MS was disconnected for less time using sequence E15 than legacy solution L-75. Using the legacy solution there were three handovers that lasted longer than $5\,s$, whereas for MROAD, the MS was disconnected for, at most, 3 s. Note that in Figure 4.7(a) when using legacy solution L-75, the MS once associated with an AP that was not on the experimental route. However, although it provided high signal strength at a given moment, it did not offer good connectivity. Finally, changes in signal strength show that MROAD experienced higher signal strengths from the beginning of the experiment until around time 70 s when there were two long periods where the RSS dropped below -80 dBm. Note that the amount of data received significantly decreased when the received signal dropped below -80 dBm. These observations confirm the results shown in Figure 4.2(a) and the study conducted in [Castignani 2012a].

Figure 4.8 shows the mean data rate of the intensive UDP flow and the corresponding associated time with a 99 % confidence interval for each scenario. The mean data rate is computed as the ratio of the amount of data downloaded over a complete run and the overall experimental time. We consider the associated time to be the time from when the MS associates with an AP until the beginning of the scan that precedes the next handover. With a maximum data rate of 1.5 $Mbps$, legacy solutions provided lower data rates than MROAD, which reached up to 1.8 $Mbps$ (bandwidth limited to 2 $Mbps$). We observe that the MS was associated up to 90 % of the time with MROAD (E15), compared to up to 72 % with the legacy solutions (L-75). Note here that the associated time and the mean data rate follow the same trend. We can explain this by the fact that the server was continuously sending packets such that when the link between the MS and the AP was up, the MS would immediately receive packets. However, this assumption is not verified for E09 to E11. An explanation of this anomaly is proposed at the end of this section.

Next, we investigate disconnection duration and signal strength distribution. Figure 4.10 illustrates the distribution of the disconnection duration for all experimental runs. The disconnection time of the legacy solutions shows a bimodal distribution. The first mode was between 0.7 s and 1 s, which includes 48 % of cases in L-70. The second mode was between 7 s and 8 s, with a probability of 13 %. In contrast, a disassociation time of 0.2 s was most frequently observed with MROAD. Two factors may explain this. First, legacy solutions performed a full active scan, i.e., they looked for APs on all channels before initiating the handover to the next AP. This scan was costly in terms of time and caused long disconnections. MROAD, on the other hand, only needed to probe an AP that already been selected by the AP selection algorithm. The difference between the disconnection experienced with MROAD and the legacy solutions suggests that MROAD probes of the selected AP do not have a significant impact on disconnection duration. Second, MROAD always chooses the most appropriate location to perform the handover. This is unlike legacy solutions, which triggered a handover based only on the RSS of the current AP, and sometimes chose APs that were not on the experimental route, leading to poor connectivity (see Figure 4.7(a)). Finally, Figure 4.9 shows the distribution of the signal strength received by the MS and the average signal strength recorded over all experiments for each scenario. Although the median signal strength of the E15 and legacy solutions were similar, the legacy solutions had a higher signal strength on average. This is because in the legacy handover approach the MS remains disconnected for longer periods (compared to MROAD), which prevents it from sensing very low RSS (see Figure 4.7(a) and Figure 4.7(c)).

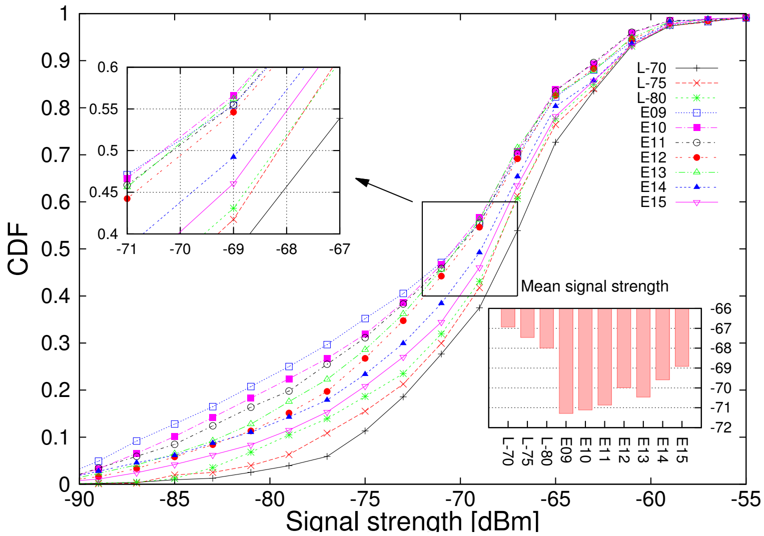

Figure 4.9: Signal strength distribution

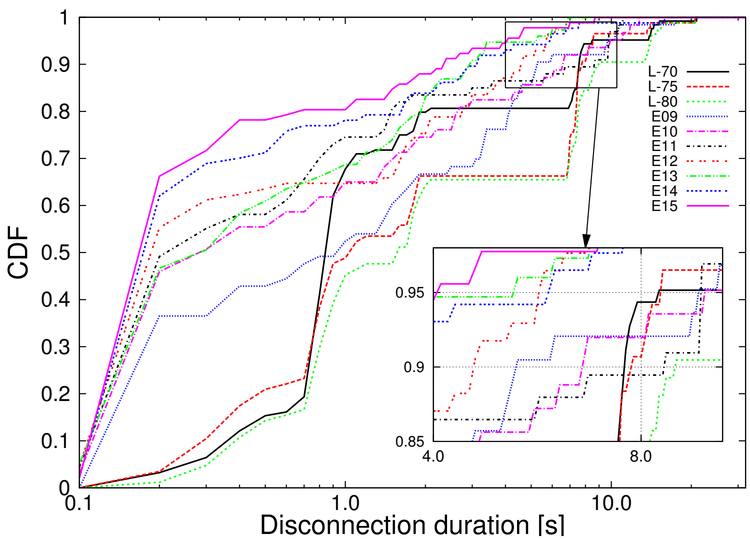

Figure 4.10: Disconnection duration distribution

Handover sequence comparison This section focuses on the handover sequences selected using the solutions obtained by the bi-objective optimization problem described in Section 3.3.3. Each solution of this problem is a handover sequence with a length varying between 9 (for E09) and 15 (for E15). Consequently, this study aimed to find the best trade-off between decreasing the number of handovers and maintaining good signal quality (even if this implies more disconnections).

First, we compare two samples of a representative run of the shortest and the longest sequence for MROAD (Figure 4.7(b) and Figure 4.7(c)). The downloaded data rate (red line) increased more rapidly for the E15 sequence than for the E09 sequence. The data rate of E09 experienced four plateaus vs. three for E15. Moreover, E09 plateaus lasted between 5 *s* and 10 *s* whereas those of E15 lasted between 2 s and 5 s. In the first 95 s of the experiment, the MS performed 7 handovers for E15 for a mean connection time of 13.6 *s* per AP, while for E09 it performed 5 handovers in 205 s for a mean connection time of 21 *s* per AP. The fact that the MS stayed connected longer to the same AP had an impact on the RSS, which dropped below −80 dBm for 20 s for E09 compared to 5 s for E15.

Figure 4.8 shows that the associated time and the download data rate increase as the length of the handover sequence increases. In particular, we observed two possible scenarios: those with best performance (from E12 to E15); and those offering worst performance (from E09 to E11). The data rate and relative association time of the first group followed the same trend, whereas in the second group the data rate significantly decreased as the sequence length decreased even if they remained associated for approximately the same length of time. Although Figure 4.9 shows similar mean signal strengths for scenarios E09 to E13, the distribution of these strengths reveals a clue about the difference between E11 and E12. The distribution of the lowest signal strengths for the group providing the worst performance was higher than for the other scenarios. The representative samples of changes in RSS and download data rate shown in Figure 4.7 highlights that such low signal strengths were not able to provide good connectivity at L3. This is because the packets error rate is high under these conditions. Although the number of disconnections decreased when the sequence length decreased, Figure 4.10 confirms that the shortest sequences have longer overall disconnection times. In particular, note that there is the same division into two groups. In scenarios E12 to E15, the MS was not disconnected for longer than 8 *s*, compared to scenarios E09 to E11, where 7 % and 10.5 % of disconnections were longer than 8 *s* (see zoom in Figure 4.10).

Note that these results are specific to the network studied because they depend on network density. In fact, municipal hotspot networks are usually sparse because the range of an outdoor AP is large (up to 300 m) and they are expensive to deploy. In contrast, CNs, which are usually very dense, would lead to a different trade-off favoring a lower number of handovers as the coexistence of multiple APs providing similar RSS would lead to the ping-pong effect and reduce connection stability.

4.4.2 COPER

4.4.2.1 Experimental setup

As for earlier studies, the experiment presented here were performed in the Gasperich neighborhood of Luxembourg. However, the route was 2.6 km long and comprised 18 roadside APs. Road furniture included one traffic light, multiple stop signs, and various intersections. However, as the speed limit was low (i.e., 30 km/h) for almost all the route, the duration of each individual experiment was roughly constant. The average duration of a single experiment was 353 s with a standard deviation of 14 s.

In order to provide meaningful results, the performance of COPER is compared to other solutions: (i) a network manager[4] (netMan), (ii) a modified version of the WPA-supplicant with an RSS threshold of -75 dB, and (iii) MROAD. We carried out this comparative study of netMan and the WPA-supplicant in order to evaluate how legacy scanning-based approaches perform compared to COPER and MROAD, which are specifically designed for vehicular communications. The comparison of COPER and MROAD examines the potential impact of the direction detection process on MS communication. The experiments are performed as follows. The MS obtains an IP address through DHCP and retains it for the duration of the experiment, as the HotCity network supports L2 handovers. When the vehicle is at the departure point, the MS initializes a 2.5 Mbps downlink UDP flow from a remote server in our laboratory. During the experiment, we captured received packets using *tcpdump* and recorded the Linux kernel log, the GPS location of the vehicle, and network interface statistics (e.g., RSS of the current AP, amount of data received and transmitted, packet loss). The experiments were repeated five times for each potential solution in order to obtain representative averaged values.

4.4.2.2 Evaluation of AP selection

Overall % of RS	Median value (dB)	Relative % of RS where $\Delta S < 5$dB	Relative % of RS $RSS < -80$ dBm
19.3	2.6	70.1	2.8

Table 4.2: Distribution of ΔS when $\Delta S > 0$

In this section, we focus on the COPER AP selection process. As previously mentioned, AP selection is suboptimal when the handover decision is delayed until the next predicted direction is taken, or when the handover is misplaced due to an incorrect RSS estimate. The spatial shifts caused by these two issues are analyzed in Figure 4.11, which shows that the handover location shifts due to a lack of information about the

[4] https://wiki.gnome.org/Projects/NetworkManager

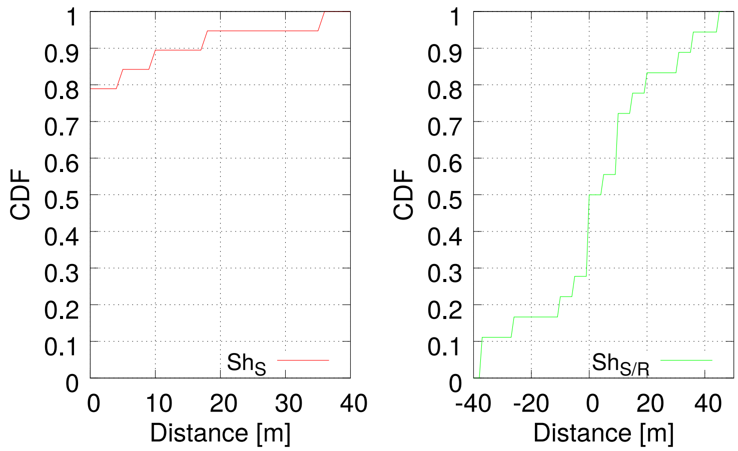

Figure 4.11: CDF of Sh_S and $Sh_{S/R}$

vehicle's direction (Sh_S) are much shorter than those caused by an erroneous RSS estimate ($Sh_{S/R}$). Sh_S is equal to zero in almost 79% of cases compared to 21% for $Sh_{S/R}$. Moreover, Sh_S is, in general, below 20 m whereas $Sh_{S/R}$ is uniformly distributed in the range $[-40m; 40m]$. We conclude that handover misplacements are mainly due to the difference between the simulated RSS (stored in the CDB) and the real RSS. The distribution of RSS differences between the current and best AP (see Table 4.2) reveals that the RSS of the current AP is the best available in more than 80% of RSSs on the run. Therefore, the misplacement of the handover has a negligible impact on RSS. Indeed, the median RSS difference relative to the best AP is only 2.8 dB and in 70.1% of the cases is under 5 dB. In addition, there are not many cases where suboptimal AP selection implies a critically low RSS (i.e., below -80 dBm): we observed $\Delta S > 0$ in only 2.8% of RS. As a result, we conclude that CORNER modeled RSSs are realistic enough to be used for AP selection.

4.4.2.3 Connection performance of the MS

In this section we analyze several metrics related to MS connectivity: the data rate of the downlink UDP flow at the MS, MS disconnection time, and the RSS of the current AP.

Network performance comparison Figure 4.12 shows the distribution of the mean data rate of the downlink UDP flow. This shows that the total time during which the MS does not receive any data from the network is much longer for both netMan and the WPA-supplicant-based handover (around 20%) than for MROAD and COPER (less than 1%). Table 4.3 shows that MROAD and COPER have both high association time

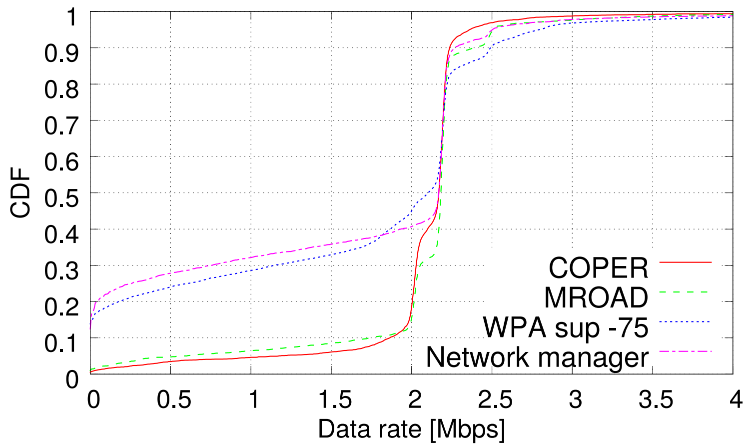

Figure 4.12: CDF of the data rate

	Estimated PDR	Associated time	Number of AP association per run	Distinct AP associated (outside roadside)	Mean connected time
NetMan	0.79 ± 0.05	88.2 ± 0.68	15.4 ± 1.17	15 (3)	23.79 ± 13.25
WPA sup -75	0.80 ± 0.02	78.8 ± 0.47	39.7 ± 3.12	25 (6)	7.15 ± 0.87
MROAD	0.97 ± 0.03	99.06 ± 0.17	19 ± 0.00	18 (0)	15.25 ± 1.43
COPER	0.95 ± 0.02	99.08± 0.14	19 ± 0.00	18 (0)	17.17 ± 1.33

†Corresponding to a confidence level of 95% assuming a normal distribution

Table 4.3: Comparative performance of handover solutions with confidence intervals†.

(above 99%) and high PDR (above 95%). Similarly, netMan and the WPA-supplicant have approximatively the same PDR (around 80%). Note that the PDR is an estimate, as in our experiments the server the MS connects to sends a flow of 2.5 Mbps, which is higher than the capacity of the network.

Figure 4.13 shows the distribution of the disconnection time for all approaches. The median value is around 1200 ms (maximum 5385 ms) for the WPA-supplicant and 1300 ms (maximum 7684 ms) for netMan, versus around 100 ms (maximum 2277 ms and 1325 ms) for MROAD and COPER respectively. This is explained by the fact that the WPA-supplicant and netMan trigger AP scanning at each handover, which could also partially explain why the MS receives no data for around 20% of the time using these solutions (see Figure 4.12). We conclude that avoiding the scan significantly reduces disconnection periods.

Furthermore, the RSS distribution (Figure 4.14) again shows a large difference

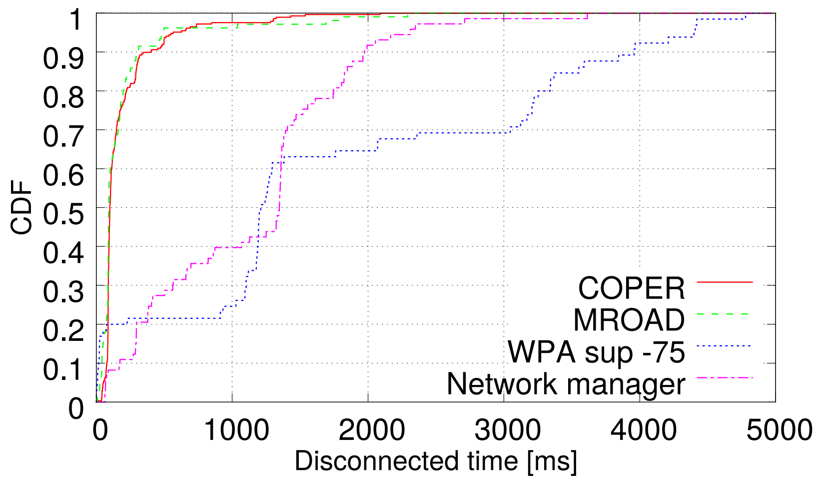

Figure 4.13: CDF of the disconnected time

between the two sets of solutions. RSS was under −80 dBm in 21% and 33% of cases for the WPA-supplicant and netMan respectively, while such a low signal level was only observed in 3.3% and 2.8% of cases for MROAD and COPER respectively. Moreover, the median observed RSS is around −64 dBm for MROAD and COPER versus −67 dBm for the WPA-supplicant and −71 dBm for netMan. Again, the RSS distributions for MROAD and COPER are similar and close to the distribution of the maximum RSS sensed over the experimental path (Figure 4.3). Figure 4.2(a) shows that when the RSS drops below −80 dBm, the data rate decreases significantly.

In contrast, the distributions for netMan and the WPA-supplicant differ significantly. Although NetMan is more affected by low RSS than the WPA-supplicant, this seems to have no impact on the data rate nor on the PDR. The RSS distribution for values greater than −65 dBm for COPER and MROAD are not equal. This could be because AP selection for COPER can be delayed compared to MROAD (cf. Section 4.4.2.2). Nevertheless, this has no significant impact on the data rate because, due to ISP bandwidth capping, all RSS above −77 dBm provide equivalent data rates (cf Figure 4.2(a)).

Next we consider association time, the number of successful associations, and the mean association time (see Table 4.3). Association time starts when the MS associates with an AP and ends at the beginning of the scan that precedes the next handover. These results show that, using MROAD and COPER, the MS is almost always associated with an AP (99% of the tune), while this is the case for only 88.2% and 78.8% of the time with netMan and the WPA-supplicant respectively. We observe a significant difference in mean connected time with these two solutions. The MS remains connected with an AP for an average of 23.79 s for NetMan and 7.15 s for the WPA-

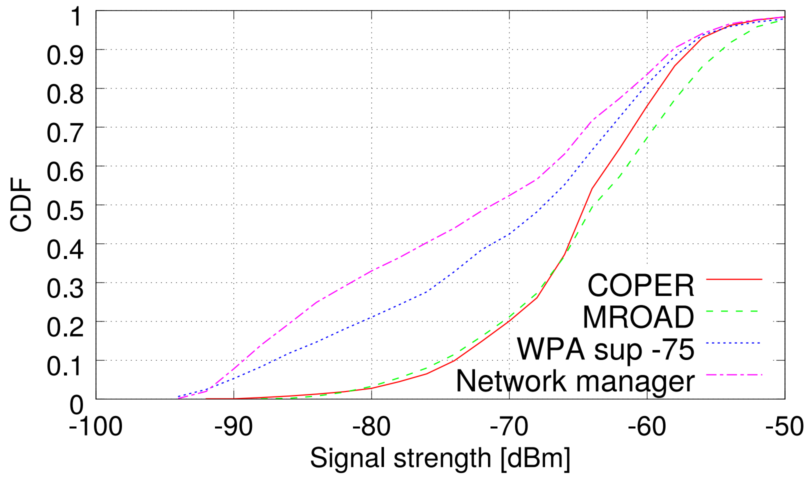

Figure 4.14: CDF of the sensed RSS

supplicant versus 15.25 s and 17.17 s for MROAD and COPER respectively. In the case of the WPA-supplicant, the higher re-connection frequency leads to shorter connection duration, which reduces the stability of the MS connection - unlike netMan that is even more durably connected than MROAD and COPER. Figure 4.13 shows that the WPA-supplicant causes long disconnections due to scanning, thus, this solution is more impacted by disconnections than netMan because it triggers handovers much more often. Note that netMan and the WPA-supplicant are affected by different factors: poor link quality during long periods for the first; and frequent disconnections for the second. These two factors seem to have a similar impact, as both have the same data rate distribution and PDR. These results also show that MROAD and COPER offer long and stable associations (the sequence of selected APs remains the same over the different runs).

We note that AP selection becomes unstable when it is only based on instantaneous scanning results. Both netMan and the WPA-supplicant select between three and eight APs that are not on the vehicle's route; these APS are unable to provide a sustainable RSS. Also, roadside APs are not always selected in their order of appearance. In contrast, MROAD and COPER only selected the 18 APs that provided the best signal all along the path.

Note that COPER slightly outperforms MROAD in terms of disconnection time. Three factors may explain this difference. The first is handover spatial shifts caused by the lack of information about the vehicle's direction (highlighted in the previous section). It may also be due to the fact that the experimental conditions were not exactly the same when evaluating MROAD and COPER, and because the performance of a given solution can vary slightly from one run to another. Nevertheless, we can conclude that COPER

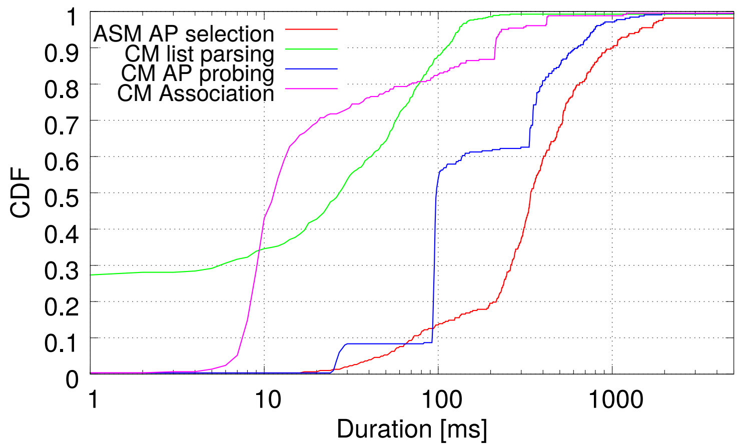

Figure 4.15: CDF of the handover phase duration

performs as well as MROAD with the advantage that it requires no prior knowledge of the vehicle's path, which greatly increases flexibility.

Impact of AP selection and the handover process Finally, we study the impact of the AP transition process on MS connectivity. We distinguish between the AP selection phase, which involves the offline computation to find the best AP, and the execution of the handover, where the CM parses the handover list, probes the selected AP and associates to it. Figure 4.15 illustrates the duration of the different elements that make up these two phases.

The need for offline computation implies that AP selection is performed by the ASM and the candidate AP list is parsed by the CM before a handover is attempted. The median duration of ASM AP selection and candidate AP list parsing are 340 ms and 27 ms respectively. This is a long time compared to the time taken to execute the handover, which takes a median of 98 ms for AP probing and 12 ms for authentication and association. However, as AP selection is performed offline, its median duration (under 400 ms) only reflects the reaction time of COPER when a new RP is discovered. We can, thus, conclude that even in the worst case, AP selection only slightly delays handover.

During the probing phase in COPER, the MS sends a broadcast Probe Request on the channel where the selected AP is operating and waits for Probe Responses from all APs in this channel for a fixed time. On the other hand, in the standard handover, the MS repeats the same process on all channels. This results in a much shorter median duration for COPER (98 ms) compared to the standard handover used by netMan (1143 ms). Note that the AP probing duration distribution has three modes that correspond to: (i) the case where the MS retrieves the selected AP from recently-stored scan results (∼20 ms); (ii) the MS received a Probe Response from the selected AP after the first attempt (∼100 ms); and (iii) the MS needs several attempts to probe the selected AP (longer than 200 ms).

The handover phase of COPER uses probing and association methods available in the existing Linux kernel module, which could be optimized. For instance, probing could be improved by specifically targeting the Service Set ID (SSID) of the selected AP, sending multiple probes within a short time interval or timing out more quickly while waiting for Probe Responses.

4.5 Discussion

In this section we discuss the performance of the proposed solutions with respect to the two objectives proposed in this thesis.

Using a second interface In Section 4.3 we showed that using a second interface with the dynamic algorithm significantly reduces handover latency and could even eliminate it. In addition, although beyond the scope of this thesis, this solution could also be used to reduce handover latency if the transition is performed at the L3. The make-before-break approach potentially allows a mobile user to experience seamless handovers though networks other than the hotspot network studied here (in particular CNs). As a result, the approach is very effective with regard to the first objective of this thesis. However, as we showed in Section 3.2.3, the AP selection process considered here relies on scanning and is vulnerable to the detection of undesirable APs and sudden disconnections due to the loss of the LoS. A second interface only provides a partial response to this issue. Consequently, it needs to be implemented alongside another solution intended to optimize AP selection (like MROAD and COPER) in order to fulfill the requirement.

MROAD and COPER Both MROAD and COPER rely on the principle of removing scanning and replacing it with prior knowledge of road topology, the location of APs, and RSS simulations. By removing the need to scan, these solutions significantly decrease handover latency compared to legacy scanning-based solutions (from more than 1 s to around 100 ms). Although improvements to the driver may reduce this latency even further, it is impossible to remove it completely due to candidate AP probing and the IEEE 802.11 standard association process. As a result, the make-before-break approach performs better than MROAD and COPER in reducing handover latency. In contrast, AP selection was improved by introducing prior knowledge of the network and road topology. Consequently, MROAD and COPER successfully address the second objective of this thesis (see Section 4.4.2.2). We also highlighted that the impact of handover location shifts due to the lack of information about the vehicle's direction is negligible. Nevertheless, AP selection could still be improved by introducing more dynamic network information so that candidate APs are chosen with respect to their actual state (including load balancing, temporary sources of interference, etc.). This concern is an element of the discussion presented in the next chapter.

In conclusion, both MROAD and COPER, which were designed to meet the objectives of handover latency reduction and improved AP selection lead to a significant improvement in vehicular communications through IEEE 802.11 networks (99% L2 connectivity and a PDR of 95%). Although similar network performance was achieved with the make-before-break approach, connections are less stable and cannot anticipate future network conditions. In contrast, the approaches proposed in this thesis offer short-term prediction of network conditions that make it possible, among other advantages, to predict white areas and anticipate connection loss by switching to another network.

Chapter 5

Conclusion and perspectives

5.1 Concluding remarks

Currently, a large part of mobile communications are provided by cellular networks (UMTS, HSPA, LTE, and in the near future LTE-A). These wide area networks provide high throughput and support high-speed mobility. At the same time, IEEE 802.11 network density has grown significantly over the past decade, particularly in urban areas. Although their deployment is unsystematic, they are able to provide high throughput in almost all public urban areas. In order to obtain maximum benefit from these heterogeneous networks, an MS must be able to make efficient connections. However, as handovers lead to frequent disconnections, the optimization of mobile communications though IEEE 802.11 networks remains challenging.

This thesis focuses on the optimization of vehicular communications through IEEE 802.11 networks. As handover latency is fixed, the impact of the handover increases with speed, which has a major impact on vehicular communications. In addition, vehicular mobility creates large variations in RSS due to high speeds and urban topology. The analysis of the legacy IEEE 802.11 handover under vehicular mobility showed that the discovery phase, i.e. scanning, is the main contributor to latency. Furthermore, it highlighted the inability of scanning to provide accurate information on nearby APs, leading to an inefficient AP selection policy. We therefore identified two objectives intended to reduce the impact of the handover, namely: (i) reduced handover latency, and (ii) improved AP selection that takes into account short-term changes in the network. The first objective aims to make handovers smoother, while the second changes the handover paradigm from a negative statement (avoid imminent disconnection) to a positive one (select a candidate AP as soon as it becomes the best AP). This change of paradigm corresponds to the *Always Best Connected* concept proposed by Gustafsson and Jonsson [Gustafsson 2003]. This paradigm aims not only to avoid sudden disconnections due to vehicular mobility, but also to optimize RSS, and thus, received throughput.

The first approach used a legacy scanning-based handover technique with the introduction of a second interface to connect to an IEEE 802.11 hotspot urban network consistent with the make-before-break principle. This purpose of this approach (which has been tested in the literature but only in indoor scenarios) was to study the negative impact of scanning-based AP selection and show that it was possible for an MS to be seamlessly connected to IEEE 802.11 networks. The results of experiments performed in a residential area highlighted the instability of AP selection (10 different AP sequences were selected in 10 experiments), while scanning results were also inconsistent. In addition, a naive, scanning-based approach often selected APs that were not on the vehicle's route and which were not able to provide good RSS over long periods due to the rapid loss of the LoS. Although AP selection was unstable, using a second interface to anticipate the new connection allowed the MS to remain connected on average 98% of the time with a PDR of 95%. Disconnections were mainly due to handovers that were triggered too close together in time. These results suggest that the make-before-break approach is very efficient with regard to the first objective of this thesis. Furthermore, although it was not shown here, it is likely that it can also reduce (or even eliminate) L3 handover latency. However, it does not address the second objective of this thesis.

Unlike the make-before-break approach, the main goal of the second approach was to optimize AP selection. This consists in providing prior knowledge of the network topology, the location of APs, and their simulated RSS to the MS (rather than scanning to obtain results). This allows the MS to select APs based on an accurate prediction of changes in the network. Consequently, the MS can select the AP providing the highest RSS for a given period, and identify the best location to perform the handover. This was determined as the location where the RSS of both current and candidate APs are simultaneously maximized. The first proposed solution was MROAD. MROAD assumes that the vehicle's route is known on departure, which makes the prior selection of APs possible. Handovers are triggered based only on the current location of the vehicle, as the MS knows about candidate APs in advance and the best handover location. The analysis of the various possible AP sequences on the experimental route highlighted a bi-objective optimization problem; between RSS maximization on the one hand, and the minimization of handovers (implying disconnections) on the other. We identified the set of sequences making up the Pareto optimal front and tested them. By avoiding scanning latency, disconnections were reduced to such an extent that priority could be given to the solution that maximised RSS. However, this result is only useful in a low-density network such as the hotspots considered in our experiments. The second contribution was COPER. This solution relies on the same principles as MROAD except that it does not require prior knowledge of the vehicle's route. Instead, it implements a direction detection module that analyses the vehicle's trajectory in real time. It assumes that the vehicle trajectory can often be predicted as it follows that road layout, while uncertainty only occurs at an intersection. Field experiments concluded that the results from COPER and MROAD were similar, although with COPER vehicle direction was detected in real time, connection time was 99%, and PDR was over 95%.

5.2 Future work

The approach proposed here has been evaluated in a controlled environment, namely a city-wide IEEE 802.11 hotspot network characterized by uniform AP density and Layer-2 (L2) handover support. This environment makes it easy to gather all the information needed to model the network, and as the network topology does not change frequently, the CDB remains reliable over time. In this section we discuss potential enhancements to COPER to improve performance, extend it to other Internet access technologies and scale it up.

5.2.1 Using COPER with other IEEE 802.11 networks

5.2.1.1 IEEE 802.11p networks

IEEE 802.11p networks have been specifically designed for vehicular communications. Consequently, it is important to evaluate COPER using these networks. However, as IEEE 802.11p has not yet been deployed it is not possible to implement field tests. In addition, to the best of our knowledge there is no study that compares IEEE 802.11b/g and IEEE 802.11p outdoor deployments.

IEEE 802.11p and the IEEE 802.11b/g/n hotspot network we used in our experiments differ in terms of frequency (2.4 GHz in our experiment versus 5.9 GHz for IEEE 802.11p), channel bandwidth (20 MHz in our experiment versus 10 MHz for IEEE 802.11p), and maximum transmission power [802 2010]. Nevertheless, both hotspot IEEE 802.11b/g/n deployments used in the evaluation of COPER and IEEE 802.11p deployments consist of a fixed set of roadside APs.

An evaluation of an experimental IEEE 802.11p network in a urban area is presented in [Gozálvez 2012]. The study evaluates IEEE 802.11p RSUs in various urban mobility settings (non LoS, bridges, elevation, trees, etc.). It shows that IEEE 802.11p communication capability is significantly reduced when the vehicle leaves the LoS of the AP. In *strong Non Line-of-Sight (NLoS) conditions* i.e., due to sharp curves or the presence of buildings, connectivity is lost within a few meters of leaving the LoS. IEEE 802.11p communications are also very sensitive to variations in elevation and the presence of vegetation or heavy vehicles. Despite these limitations, the authors showed that IEEE 802.11p networks can provide reliable connectivity (packet delivery ratio > 0.7) to a mobile vehicle at distances ranging from a little under 100 m to a little more than 750 m. We can, thus, conclude that it is feasible to deploy city-wide IEEE 802.11p networks that can provide good signal strength all along the route taken by a given vehicle (or at least, a large part of it).

In this case, the results presented in [Gozálvez 2012] suggest that a scenario based on a city-wide IEEE 802.11p network may lead to results that are comparable to those highlighted in the previous section. However, no definite conclusions can be drawn in the absence of field tests.

In addition, IEEE 802.11p differs from other IEEE 802.11 networks as it does not define an association procedure between the MS and the AP. Furthermore, the IEEE 802.11p WAVE extension [IEE 2011] and ITS G5 [ETSI 2010] permit the use of a second radio dedicated to communication on the Service Channel (SCH) that can be used to scan nearby RSUs (amongst other things). This completely removes the impact of the handover on vehicular communications, which was one of the goals of COPER. However, the scanning-based AP selection approach that is proposed in IEEE 802.11p is potentially inefficient (as shown here for IEEE 802.11 networks). As a result, COPER could also be used with IEEE 802.11p networks in order to improve AP selection by taking into account vehicular mobility.

5.2.1.2 Community networks

Residential and community networks provide households with Internet access and are characterized by their unsystematic deployment. However, the ever-growing number of community networks offers widespread coverage in urban and suburban environments, which could be used by COPER. In this section we discuss the potential gains and the issues raised when using IEEE 802.11 community networks with COPER.

Figure 5.1(a) shows some interesting results. First, RSS is more uniformly distributed in the residential AP network than the HotCity network. Second, areas where the RSS is lowest (shown in orange) are wider for residential APs than for HotCity APs. In general, HotCity APs provide higher RSS than residential APs (between -84 dBm and -50 dBm versus -88 dBm to -56.5 dBm). Nevertheless, we note that even in some residential areas RSS is high(shown in green). This phenomenon occurs in the most populated areas (the large, inhabited buildings shown at top of the figure versus greater numbers of individual houses in other areas). Finally, it is important to note that both networks cover the entire experimental path.

Regarding AP selection, the high density of APs in community networks in certain areas implies that the AP selection process designed for lower-density networks must be adapted. In a context of high AP density, the trade-off between minimizing the number of handovers, and maximizing their quality (studied in Chapter 3, Section 3.3.3) tends to favor handover quantity (at the expense of quality). This is because the probability of having multiple APs with similar RSS is much higher than for hotspot networks. This phenomenon leads to situations where a given AP remains the best AP only in a very small area, forcing the MS to perform too-frequent handovers. Therefore, the best AP should be defined as the AP that is able to provide good RSS over a sufficiently long period, even if it does not always provide the highest RSS.

Unlike the IEEE 802.11 hotspot network studied in this thesis, or IEEE 802.11p networks, CNs do not provide mobility support. An MS must renew its IP configu-

(a) Residential APs' highest RSS distribution

(b) HotCity APs' highest RSS distribution

Figure 5.1: Hotspot and residential networks RSS in the experimental path

ration each time there is a handover to a candidate AP that belongs to a community network. This has two main consequences: (i) the handover delay significantly increases compared to other networks, and (ii) a network mobility management protocol such as MIP is required to provide session continuity.

These delays are not consistent with the objective of eliminating the handover impact defined in this thesis, and the use of COPER in such networks would not significantly decrease handover latency. One way to overcome the problem may be to add a second interface (this solution is discussed later in this chapter).

In addition, a high number of APs and frequent changes in network deployment raise questions about information collection and dynamic updates to the Context DataBase (CDB). This is discussed later in this section.

5.2.2 Providing long-term AP selection

In the current version of COPER, the AP selection process is triggered when the vehicle's direction is detected with a high level of confidence. This constraint sometimes delays AP selection and is not error-free. Moreover, this approach provides a limited view of the network environment in the very near future.

We have already described an AP classification mechanism that aims to improve the resilience of COPER. This enables MS to immediately select a new candidate AP from the list of APs in the current RP, in the case where the previously selected one is not available. The classification is based on the handover quality parameter (see Chapter 3, Section 3.4.4). However, this mechanism only provides partial protection against the unavailability of a selected AP (if, for example, it is switched off or bandwidth is insufficient) as there may be no other AP in the current RP. In this case, the MS must wait for the next direction prediction in order to select an alternative AP.

There are several contributions in the literature that use large datasets to perform mobility prediction and provide long-term AP selection. Zhang et al. [Zhang 2014] propose a cloud-based prediction system based on the mobility history of cellular network users. This approach consists in detecting periodicity in the mobility pattern using the Kullback-Leibler divergence (KLD) as a metric, and evaluating social interplay in order to identify pairs of calls that are co-located. Abu-Ghazaleh and Alfa [Abu-Ghazaleh 2010] model the mobility behavior of users in an IEEE 802.11 network as a Markov Renewal Process (MRP). This model predicts handovers and the length of stay of a user based on their prior location immediately before the transition to the current location. In [Wanalertlak 2011], the authors propose behavior-based mobility prediction using as input the location of the user, the group it belongs to (students, graduate students, faculty staff, etc.), the time of day and connection duration. AP selection, represented by an order-k Markov process, is performed by a server upon request. A similar approach is proposed in [Chen 2013] using the mobility traces of

16000 taxis to perform mobility prediction using an order-k Markov process.

The proposed approaches need to be analyzed and compared in order to identify which one provides the most accurate results. Beyond long-term mobility prediction, the AP selection can be improved by including other contextual information (discussed in next section).

5.2.3 Providing COPER with dynamic network knowledge

5.2.3.1 Towards a collaborative approach

COPER uses a Context DataBase (CDB) containing the modeled RSSs of potential APs. This knowledge reflects the RSS of the entire network coverage, but only takes into account communication between the MS and the AP.

However, in the urban environment there are multiple temporary or permanent sources of interference that can alter the RSS of the MS. First, the 2.4 GHz band is often crowded due to the high density of APs in urban areas, as shown in Section 2.1.4. Next, the urban topology features many obstacles that can impact communication between the MS and the AP. The results presented in [Gozálvez 2012] show that IEEE 802.11p communications are significantly affected by loss of LoS and elevation (see Section 5.2.1.1). This study also shows that many of the sources of interference are not static. The two main dynamic sources cited by the author are vegetation (which changes according to the season) and heavy vehicles (depending on road traffic). Note that communications in the 2.4 GHz band are affected in the same way.

In addition, the presence of an AP in the CDB does not guarantee that this AP can actually be used by the MS. It could, for example, be switched off or overloaded and unable to provide sufficient bandwidth. In order to improve the AP selection of COPER, the process should take into account changes in the network in question. This raises two challenges related to: (i) the collection of information, and (ii) the storage of this information in the CDB.

<table>
<tr><th colspan="2">Information</th><th>Frequency of change</th><th>Durability</th><th>How to get this information</th></tr>
<tr><td rowspan="2">Change in network topology</td><td>New AP</td><td rowspan="2">Not frequent</td><td>Long term</td><td rowspan="2">ISP & auto detection</td></tr>
<tr><td>AP removal</td><td>Mid / Long term</td></tr>
<tr><td colspan="2">AP load / Bandwidth allocation</td><td>Very frequent</td><td>Short term</td><td>IEEE 802.11k & RSVP</td></tr>
<tr><td rowspan="2">Dynamic sources of interferences</td><td>Trees & Vegetation</td><td>Not frequent</td><td>Long term</td><td>Model</td></tr>
<tr><td>Mobile obstacles</td><td>Very frequent</td><td>Short term</td><td>/</td></tr>
</table>

Table 5.1: Dynamic information about APs

Collecting dynamic information The network information that can be used by the MS, described in Table 5.1, should be categorized according to the way the MS can retrieve it. The first type of information is change in the network or road topology that can be either positive (the implementation of a new AP or the construction of a new road), negative (the removal of an AP, or a road closure) or neutral (a change in AP configuration or modifications to a road). Such events are usually infrequent and last for a long time, except in the case where an AP is temporary switched off (for instance during construction work). Only the entity that maintains the network or the road, i.e. the ISP or the municipality, are aware of these topology changes. However, it is possible for an MS to detect a new AP if it continues to scans, so longer as it not is a matter of AP selection, but of AP discovery (for instance when the interface is idle). In the same manner, AP removal can also be detected after a number of failed connections by different MSs, although this is not the case for road topology changes. The second type of information is AP load and the bandwidth allocated by it. Depending on the time of day, this setting can vary widely, and frequently. This information can be retrieved from the AP itself using the IEEE 802.11k protocol, which provides a site report describing the use rate of nearby APs. Resource reSerVation Protocol (RSVP) can be used to allocate a given bandwidth to the selected AP before the transition to ensure satisfactory QoS. Finally, the third type of information is dynamic sources of interferences (highlighted in [Gozálvez 2012]). Here, a distinction has to be made between the two main sources; vegetation and heavy mobile obstacles. In the first case, variation is often seasonal and thus predictable. As a result, it can be taken into account while modeling the network. On the other hand, the impact of mobile obstacles is sudden and difficult to predict. It appears infeasible to collect and include such events in the CDB. Nevertheless, the probable occurrence of these events should be taken into account by COPER. The selection of the AP that provides the best signal is a way to mitigate the impact of such events, as the probability of communication disruption is higher if the signal is low. If the impact on the RSS is too high, and if the MS is multihomed, a resilience mechanism is needed to trigger a fast-switch to another access network.

Dynamic information storage The dynamic storage of information in a database does not present any technical problem. However, modification of the network or road topology require the re-modeling of the RSS in the area in question. The requirement for the MS to create a new model upon receipt of an update to the network or road topology adds a significant amount of complexity to COPER. In addition, such a computation is not always feasible in an embedded device. Furthermore, allowing the MS to update the CDB would lead to a situation in which the CDB shared by several MSs would diverge depending on the information provided by each station. Finally, the content of the CDB would quickly become out-of-date.

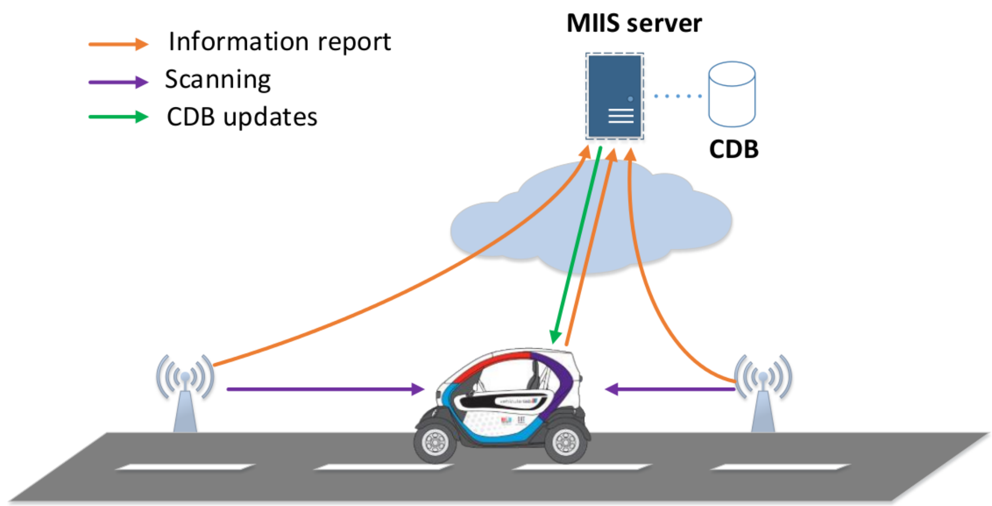

Figure 5.2: Information exchange in the proposed IEEE 802.21-based solution

A collaborative approach could overcome this problem. Here, the aim is to improve the knowledge of each MS individually by sharing knowledge with other MSs. This collaboration is possible for all dynamic information that varies over the long term (network topology changes and seasonal changes in RSS). Short-term changes, such as AP load cannot be shared by MSs as it very quickly becomes out-of-date. Collaboration between MSs implies a that there is a place to share all of the information that has been gathered. Cloud computing technology could provide such a shared space. COPER would, then, turn into an Mobile CrowdSourcing (MCS) application [Ganti 2011]. In addition to updating the CDB, the information provided by MSs to the cloud could be used to evaluate changes in the network in order, for instance, to detect faulty APs, or for studies of user behavior based on data-mining techniques. Cloud-based technologies would also allow COPER to be used on a larger scale, as it would not be possible to store a CDB that gathers information related to a very wide area on a mobile device. Nevertheless, even if only the master version of the CDB is stored in the cloud, this raises new challenges with regard to the retrieval of a partial dataset. Therefore, the CDB structure should allow the MS to extract only the information it requires, which is limited to an area around a given location.

5.2.3.2 Using the standard IEEE 802.21 structure

In this section, we discuss the potential modification of the IEEE 802.21 MIH standard in order to integrate COPER. The IEEE 802.21 standard defines a data structure and different types of communication (event, information and control) between IEEE 802.21-enabled elements. It also describes an MIIS server that gathers and diffuses, on

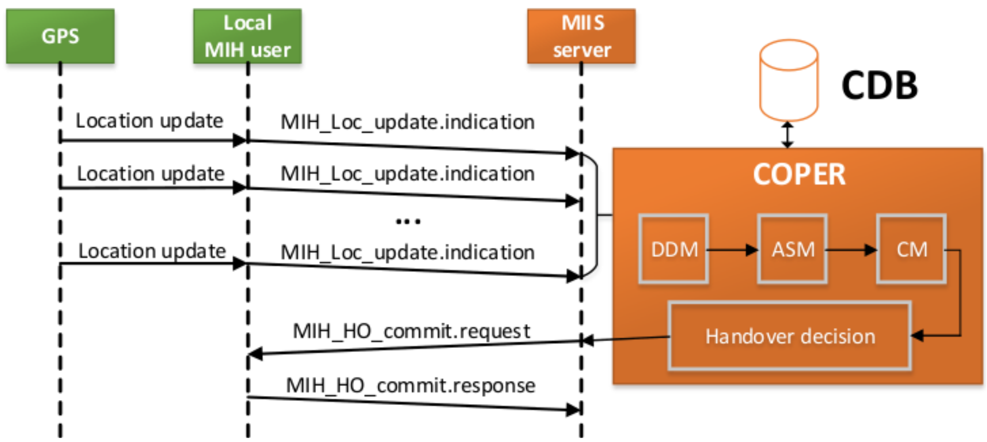

Figure 5.3: Scenario 1: Externalization of COPER

demand, useful network information.

A schme of the information exchange in the proposed solution is depicted in Figure 5.2. Consistent with the solutions already presented, we consider a wide-scale scenario where the CDB is too large to be stored locally. In addition, it is frequently updated with dynamic information. COPER could be integrated into this structure by storing the CDB in the MIIS server and/ or another COPER module. The MIIS primitives are used to feed the CDB through: (i) a scanning report sent by the MS that provides updates of the RSS along the route taken by the vehicle; and (ii) samples of the CDB that are sent to the MS depending on their location, as shown in Figure 5.3. As mentioned earlier, scanning is not used for AP selection, but in order to provide accurate information to the CDB. In order to avoid any impact on MS communications, scanning could made be passive while the MS is in use and active otherwise. Note that if APs are IEEE 802.21-enabled, they could provide real-time dynamic information about load and available bandwidth to the CDB.

Three COPER integration scenarios can be envisaged: (i) both COPER and the CDB are located in the MIIS server (see Figure 5.3); (ii) the ASM module and the CDB are located in the MIIS server (see Figure 5.4); and (iii) only the CDB is located in the MIIS server.

In the first case, direction detection and AP selection are performed by the MIIS server. The MS continuously sends GPS location updates to the MIIS server, which processes the trajectory of the vehicle and selects APs accordingly. When the vehicle reaches a handover location, the MIIS server sends a handover command to the MS, as shown in Figure 5.3. Such a solution is suitable if the MS is a very lightweight terminal, as no computation is required.

In the second case, only AP selection is performed by the MIIS server. In this case

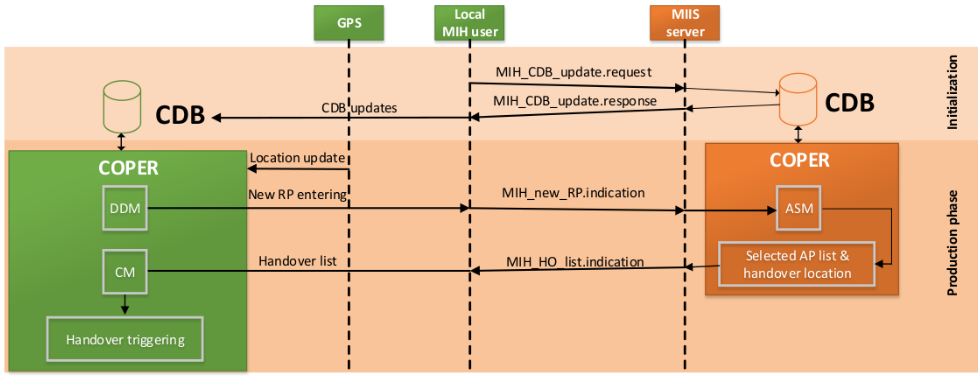

Figure 5.4: Scenario 2: Externalization of AP selection alone

the server returns a list of handover descriptions (AP and handover location). The MS detects direction and triggers the handover when the vehicle reaches the handover location. This functionality only requires the CDB to contain information about road topology and a description of APs. Communication between the MS and the MIIS server is described in Figure 5.4. At the start, the MS requests the part of the CDB that corresponds to its current location. This action is repeated when the vehicle is about to leave the area in question. When the MS has received the relevant information, the DDM starts. Upon entry into a new RP, the MIIS server is informed of this new direction and starts the AP selection process. The resulting handover list is then sent back to the MS, which is able to process it, thanks to the CM. This solution has the advantage that it requires fewer messages to be exchanged between the MS and the MIIS server, and the CDB contains less information than the current version of COPER.

In the last case, COPER is not externalized on the MIIS server. The MS only needs to request a part of the CDB, as described in the initialization phase shown in Figure 5.4. Unlike the previous case, the CDB contains the same information as in the current version of COPER.

In conclusion, extending the IEEE 802.21 standard with COPER is one way to make COPER more scalable. Currently there are three potential ways to integrate COPER into the IEEE 802.21 structure. However, further investigation is required to evaluate the different approaches in detail, in particular with regard to any delays resulting from communication between the MS and the MIIS server. With respect to the proposed architecture, new questions arise concerning the initialization of communication with the MIIS server (the MS should be able to initiate a network connection even when there is no CDB for its current location). Note also that such a modification to the IEEE 802.21 standard requires the definition of new messages consistent with the new functionality offered by this modification, such as the *MIH_Loc_update* and *MIH_CDB_update* messages shown in Figures 5.3 and 5.4.

5.2.4 Using COPER with a second interface

In Chapter 3, Section 3.2 we described a make-before-break approach where the MS uses a second interface to perform the handover to a candidate AP before disconnecting with the current AP. In Chapter 4, Section 4.3 we showed that such a solution, despite the fact that it uses scanning-based AP selection, could potentially remove handover latency, even if the handover was also performed at Layer-3 (L3) (IP configuration renewal using the DHCP protocol). On the other hand, in the current version of COPER there is still a short disconnection when the MS hands off to the selected AP (median value of 100ms). Using COPER in an encrypted network will create (incompressible) longer disconnections due to encryption key renegotiation that can only be reduced (but not eliminated) by the use of IEEE 802.11r. Therefore, using COPER with a second interface is an easy way to meet the first objective.

Furthermore, in a multipath scenario (as described in Chapter 2, Section 2.1.3), a second interface allows the MS to be connected to two APs at the same time (except during handovers) leading to an increase in send/ receive throughput.

5.3 Perspectives

In recent years, the emergence of a multitude of new wireless communication standards such as IEEE 802.11ac/ad/af/ah, in addition to the release of new frequency bands offers an opportunity to use a wider spectrum for wireless communication. The downside of this is that each new standard requires a specific radio for each available technology. Cognitive radio mitigates this problem by using a new generation of large-spectrum radios in which link layers can be quickly reprogrammed. This technology also offers new opportunities with respect to peer-to-peer communications as they can negotiate physical and link layer parameters (e.g., channel, band size, encoding technique, multiplexing mode) depending on network conditions. This technology can be used for Internet connection sharing, among other applications.

The expanding range of Internet access networks offers the potential for ubiquitous network connection. Although this is a challenging objective, here we try to show that technical solutions exist that can efficiently provide seamless Internet connections to mobile users. However, they involve an information exchange between the MS and its environment in order to allow the MS to anticipate future changes in the network. In particular, collaboration with ISPs is required in order to provide the MS with holistic knowledge of nearby networks. Like data roaming in cellular networks (which will be transparent and free in European countries from 2016), collaboration between ISPs and community networks must also be established in order to facilitate the transition between APs that belong to different CNs. The purpose of this collaboration would be to standardize and simplify user credential management in order to reduce handover latency. However, network handovers remain inevitable as each ISP must grant network access to the MS.

This raises the possibility of the emergence of a new sort of service provider that aims to always provide the best connection, through the use of multiple simultaneous Internet access points (provided by legacy ISPs). Another possibility is that the always-best-connection would become an extra service provided by existing ISPs. In any case, a cross-layer approach should be implemented in order to provide an always-best-connected service. With respect to layer-1 (L1), cognitive radio allows the dynamic selection of the spectrum according to needs in terms of coverage, bandwidth, and nearby networks. With respect to the L2, mobility is handled by protocols such as IEEE 802.21 MIH. The proposals put forward in this thesis are part of this L2 mobility management framework. With respect to the upper layer, network mobility protocols are used to guarantee session continuity, while session transfer and reachability are ensured by application layer solutions.

One potential use case, that may be considered as a priority, is providing seamless Internet access to emergency services such as ambulances, fire-fighters or the police. This would allow the exchange of critical information between personnel located at the site of the intervention and others in a command center (e.g., a hospital or police headquarters). This would make it possible to evaluate the gravity of the situation, prepare the materials necessary for a quick intervention, or provide advice to personnel in the field. Although this use case seems realistic in urban areas, emergency vehicles intervene in all parts of a country. In areas with a very sparse population, the quality of networks providing Internet access can be poor, leading to potential white areas. In this case, data communication continuity may be provided by satellite networks and cognitive radio that enable long-range communication through specific frequency bands.

Bibliography

[802 2010] *IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments.* IEEE Std 802.11p-2010, pages 1–51, July 2010. (Cited on page 101.)

[Abu-Ghazaleh 2010] Haitham Abu-Ghazaleh and Attahiru Sule Alfa. *Application of mobility prediction in wireless networks using markov renewal theory.* Vehicular Technology, IEEE Transactions on, vol. 59, no. 2, pages 788–802, 2010. (Cited on page 104.)

[Addepalli 2013] Sateesh Addepalli, Henning G Schulzrinne, Amandeep Singh and Gaston Ormazabal. *Heterogeneous Access: Survey and Design Considerations.* 2013. (Cited on page 16.)

[Aguayo 2004] Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd and Robert Morris. *Link-level measurements from an 802.11 b mesh network.* In ACM SIGCOMM Computer Communication Review, volume 34, pages 121–132. ACM, 2004. (Cited on pages 21, 35 and 59.)

[Annese 2011] Stefano Annese, Claudio Casetti, C Chiasserini, Nazario Di Maio, Andrea Ghittino and Massimo Reineri. *Seamless connectivity and routing in vehicular networks with infrastructure.* Selected Areas in Communications, IEEE Journal on, vol. 29, no. 3, pages 501–514, 2011. (Cited on pages 3 and 32.)

[Arcia-Moret 2014] A Arcia-Moret, L. Molina, N. Montavont, G. Castignani and A Blanc. *Access Point Discovery in 802.11 Networks.* In Network Games, Control and Optimization (NetGCooP), 2012 6th International Conference on, Nov 2014. (Cited on page 34.)

[Balasubramanian 2010] Aruna Balasubramanian, Ratul Mahajan and Arun Venkataramani. *Augmenting mobile 3G using WiFi.* In Proceedings of the 8th international conference on Mobile systems, applications, and services, pages 209–222. ACM, 2010. (Cited on pages 15, 16 and 24.)

[Bangolae 2006] Sangeetha Bangolae, Carol Bell and Emily Qi. *Performance Study of Fast BSS Transition Using IEEE 802.11R.* In Proceedings of the 2006 International Conference on Wireless Communications and Mobile Computing, IWCMC '06, pages 737–742, New York, NY, USA, 2006. ACM. (Cited on page 36.)

[Berezin 2011] Maria Eugenia Berezin, Franck Rousseau and Andrzej Duda. *Multi-channel virtual access points for seamless handoffs in IEEE 802.11 wireless networks*. In Vehicular Technology Conference (VTC Spring), 2011 IEEE 73rd, pages 1–5. IEEE, 2011. (Cited on page 34.)

[Biswas 2006] Subir Biswas, Raymond Tatchikou and Francois Dion. *Vehicle-to-vehicle wireless communication protocols for enhancing highway traffic safety*. Communications Magazine, IEEE, vol. 44, no. 1, pages 74–82, 2006. (Cited on page 9.)

[Brik 2005] Vladimir Brik, Arunesh Mishra and Suman Banerjee. *Eliminating handoff latencies in 802.11 WLANs using multiple radios: Applications, experience, and evaluation*. In Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement, pages 27–27. USENIX Association, 2005. (Cited on pages 3 and 32.)

[Bychkovsky 2006] Vladimir Bychkovsky, Bret Hull, Allen Miu, Hari Balakrishnan and Samuel Madden. *A measurement study of vehicular internet access using in situ Wi-Fi networks*. In Proceedings of the 12th annual international conference on Mobile computing and networking, pages 50–61. ACM, 2006. (Cited on pages 19, 20, 22, 24, 25 and 26.)

[Castignani 2011a] G. Castignani, A. Arcia and N. Montavont. *A Study of the Discovery Process in 802.11 Networks*. ACM SIGMOBILE Mobile Computing and Communications Review, vol. 15, no. 1, pages 25–36, 2011. (Cited on pages 21, 28, 31, 33 and 44.)

[Castignani 2011b] German Castignani, Lucien Loiseau and Nicolas Montavont. *An evaluation of IEEE 802.11 community networks deployments*. In Information Networking (ICOIN), 2011 International Conference on, pages 498–503. IEEE, 2011. (Cited on pages 19, 20, 21, 29 and 76.)

[Castignani 2012a] G. Castignani, A. Blanc, A. Lampropulos and N. Montavont. *Urban 802.11 Community Networks for Mobile Users: Current Deployments and Prospectives*. Mobile Networks and Applications, pages 1–12, 2012. (Cited on pages 1, 21, 24, 25, 44, 59, 77 and 86.)

[Castignani 2012b] G. Castignani, A. Lampropulos, A. Blanc and N. Montavont. *Wi2Me: A Mobile Sensing Platform for Wireless Heterogeneous Networks*. In Distributed Computing Systems Workshops (ICDCSW), 2012 32nd International Conference on, pages 108–113, June 2012. (Cited on pages 34 and 76.)

[Castignani 2013] German Castignani, Juan Monetti, Nicolas Montavont, Andrés Arcia-Moret, Raphaël Frank and Thomas Engel. *A study of urban IEEE 802.11* *hotspot networks: towards a community access network*. In Wireless Days (WD), 2013 IFIP, pages 1–8. IEEE, 2013. (Cited on pages 20, 22, 23, 24 and 45.)

[Chen 2013] Siyu Chen, Yong Li, Wenyu Ren, Depeng Jin and Pan Hui. *Location prediction for large scale urban vehicular mobility*. In Wireless Communications and Mobile Computing Conference (IWCMC), 2013 9th International, pages 1733–1737. IEEE, 2013. (Cited on page 104.)

[Deb 2001] Kalyanmoy Deb and Deb Kalyanmoy. Multi-Objective Optimization Using Evolutionary Algorithms. Wiley, 1st édition, June 2001. (Cited on page 60.)

[Deshpande 2009] Pralhad Deshpande, Anand Kashyap, Chul Sung and Samir R. Das. *Predictive Methods for Improved Vehicular WiFi Access*. In Proceedings of the 7th international conference on mobile systems, applications, and services, MobiSys '09, pages 263–276, New York, NY, USA, 2009. ACM. (Cited on pages 35 and 54.)

[Deshpande 2010] Pralhad Deshpande, Xiaoxiao Hou and Samir R Das. *Performance comparison of 3G and metro-scale WiFi for vehicular network access*. In Proceedings of the 10th ACM SIGCOMM conference on Internet measurement, pages 301–307. ACM, 2010. (Cited on pages 22, 24 and 44.)

[DLR] DLR. *Simulation of Urban MObility, [online] Available*. http://sumo.sourceforge.net/. (Cited on page 55.)

[Eriksson 2008] Jakob Eriksson, Hari Balakrishnan and Samuel Madden. *Cabernet: vehicular content delivery using WiFi*. In Proceedings of the 14th ACM international conference on Mobile computing and networking, pages 199–210. ACM, 2008. (Cited on pages 20, 24, 25, 26, 29 and 76.)

[ETSI 2004] EN ETSI. *EN 300 674-1: Electromagnetic compatibility and Radio spectrum Matters (ERM); Road Transport and Traffic Telematics (RTTT); Dedicated Short Range Communication (DSRC) transmission equipment (500 kbit/s / 250 kbit/s) operating in the 5,8 GHz Industrial, Scientific and Medical (ISM) band; Part 1: General characteristics and test methods for Road Side Units (RSU) and On-Board Units (OBU)*. Rapport technique, Tech. Rep, 2004. (Cited on page 12.)

[ETSI 2010] EN ETSI. *302 663: Draft V1. 3.2, ITS; Access Layer Specification for Intelligent Transport Systems operating in the 5 GHz frequency band*. Rapport technique, Tech. Rep, 2010. (Cited on page 102.)

[Farshad 2014] Arsham Farshad, Mahesh K Marina and Francisco Garcia. *Urban WiFi characterization via mobile crowdsensing.* In Network Operations and Management Symposium (NOMS), 2014 IEEE, pages 1–9. IEEE, 2014. (Cited on pages 1, 19, 20, 21 and 29.)

[FCC 2014] FCC. *15.407 General technical requirements., [online] Available.* http://www.hallikainen.org/FCC/FccRules/2013/15/407/section.pdf, 2014. (Cited on page 45.)

[Ganti 2011] Raghu K Ganti, Fan Ye and Hui Lei. *Mobile crowdsensing: current state and future challenges.* Communications Magazine, IEEE, vol. 49, no. 11, pages 32–39, 2011. (Cited on page 107.)

[Garelli 2011] Laura Garelli, Claudio Casetti, C Chiasserini and Marco Fiore. *Mobsampling: V2V communications for traffic density estimation.* In Vehicular Technology Conference (VTC Spring), 2011 IEEE 73rd, pages 1–5. IEEE, 2011. (Cited on page 9.)

[Gass 2005] Richard Gass, James Scott and Christophe Diot. *Measurements of in-motion 802.11 networking.* In Mobile Computing Systems and Applications, 2006. WMCSA'06. Proceedings. 7th IEEE Workshop on, pages 69–74. IEEE, 2005. (Cited on pages 12 and 22.)

[Gerla 2011] Mario Gerla and Leonard Kleinrock. *Vehicular networks and the future of the mobile internet.* Computer Networks, vol. 55, no. 2, pages 457–469, 2011. (Cited on pages 1, 8 and 9.)

[Giordano 2010a] E. Giordano, E. De Sena, Giovanni Pau and M. Gerla. *VERGILIUS: A Scenario Generator for VANET.* In Vehicular Technology Conference (VTC 2010-Spring), 2010 IEEE 71st, pages 1–5, 2010. (Cited on pages 55 and 65.)

[Giordano 2010b] Eugenio Giordano, Raphael Frank, Giovanni Pau and Mario Gerla. *CORNER: a realistic urban propagation model for VANET.* In Wireless On-demand Network Systems and Services (WONS), 2010 Seventh International Conference on, pages 57–60. IEEE, 2010. (Cited on pages 20, 29, 55 and 76.)

[González 2008] Víctor González, Alberto Los Santos, Carolina Pinart and Francisco Milagro. *Experimental demonstration of the viability of IEEE 802.11 b based inter-vehicle communications.* In Proceedings of the 4th International Conference on Testbeds and research infrastructures for the development of networks & communities, page 1. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2008. (Cited on page 13.)

[Gozálvez 2012] Javier Gozálvez, Miguel Sepulcre and Ramon Bauza. *IEEE 802.11 p vehicle to infrastructure communications in urban environments.* Communications Magazine, IEEE, vol. 50, no. 5, pages 176–183, 2012. (Cited on pages 101, 105 and 106.)

[Grunenberger 2010] Yan Grunenberger and Franck Rousseau. *Virtual access points for transparent mobility in wireless LANs.* In Wireless Communications and Networking Conference (WCNC), 2010 IEEE, pages 1–6. IEEE, 2010. (Cited on page 34.)

[Gustafsson 2003] Eva Gustafsson and Annika Jonsson. *Always best connected.* Wireless Communications, IEEE, vol. 10, no. 1, pages 49–55, 2003. (Cited on pages 2, 45 and 99.)

[Hadaller 2007] David Hadaller, Srinivasan Keshav, Tim Brecht and Shubham Agarwal. *Vehicular opportunistic communication under the microscope.* In Proceedings of the 5th international conference on Mobile systems, applications and services, pages 206–219. ACM, 2007. (Cited on pages 18 and 22.)

[Handley 2013] M Handley, O Bonaventure and U catholique de Louvain. *Internet Engineering Task Force (IETF) A. Ford Request for Comments: 6824 Cisco Category: Experimental C. Raiciu.* 2013. (Cited on page 16.)

[Holma 2011] Harri Holma and Antti Toskala. Lte for umts: Evolution to lte-advanced. John Wiley & Sons, 2011. (Cited on page 13.)

[IEE 2011] *IEEE Standard for Wireless Access in Vehicular Environments (WAVE)–Multi-channel Operation.* IEEE Std 1609.4-2010 (Revision of IEEE Std 1609.4-2006), pages 1–89, Feb 2011. (Cited on page 102.)

[Intelligence 2014] GSMA Intelligence. *Measuring mobile penetration, [online] Available.* https://gsmaintelligence.com/analysis/2014/05/measuring-mobile-penetration/430/, 2014. (Cited on page 10.)

[Ivov 2005] Emil Ivov and Thomas Noel. *Soft Handovers over 802.11 b with Multiple Interfaces.* In Proceedings of the 2nd International Symposium on Wireless Communication Systems (ISWCS'05), pages 549–554, 2005. (Cited on page 32.)

[Kassar 2008] Meriem Kassar, Brigitte Kervella and Guy Pujolle. *An overview of vertical handover decision strategies in heterogeneous wireless networks.* Computer Communications, vol. 31, no. 10, pages 2607–2620, 2008. (Cited on page 40.)

[Khan 2012a] Muhammad Qasim Khan and Steinar Hidle Andresen. *Poa selection in 802.11 networks using media independent information server (miis).* In Advanced Information Networking and Applications Workshops (WAINA), 2012

26th International Conference on, pages 454–459. IEEE, 2012. (Cited on page 34.)

[Khan 2012b] Muhammad Qasim Khan and Steinar Hidle Andresen. *Zero scanning time for 802.11 networks by using media independent information server (miis)*. In Advanced Information Networking and Applications (AINA), 2012 IEEE 26th International Conference on, pages 467–473. IEEE, 2012. (Cited on pages 33 and 34.)

[Kwak 2009] Daehan Kwak, Jeonghoon Mo and Moonsoo Kang. *Investigation of handoffs for IEEE 802.11 networks in vehicular environment*. In Ubiquitous and Future Networks, 2009. ICUFN 2009. First International Conference on, pages 89–94, June 2009. (Cited on page 35.)

[Liao 2006] Yong Liao and Lixin Gao. *Practical schemes for smooth MAC layer handoff in 802.11 wireless networks*. In Proceedings of the 2006 International Symposium on on World of Wireless, Mobile and Multimedia Networks, pages 181–190. IEEE Computer Society, 2006. (Cited on pages 29, 30 and 82.)

[Mahajan 2007] Ratul Mahajan, John Zahorjan and Brian Zill. *Understanding WiFi-based connectivity from moving vehicles*. In Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, pages 321–326. ACM, 2007. (Cited on page 62.)

[Media 2014] Informa Telecoms & Media. *Understanding today's smartphone user: An updated and expanded analysis of data-usage patterns in six of the world's most advanced 4G LTE markets, [online] Available*. http://www.informatandm.com/wp-content/uploads/2013/06/Mobidia-ITM-June-2013.pdf, 2014. (Cited on page 40.)

[Melia 2007] Telemaco Melia, Antonio de la Oliva, Ignacio Soto, Pablo Serrano and R Aguiar. *Network controlled handovers: challenges and possibilities*. Wireless Personal Communications, vol. 43, no. 3, pages 959–974, 2007. (Cited on page 40.)

[Mhatre 2006] Vivek Mhatre and Konstantina Papagiannaki. *Using Smart Triggers for Improved User Performance in 802.11 Wireless Networks*. In Proceedings of the 4th international conference on Mobile systems, applications and services, pages 246–259. ACM, 2006. (Cited on pages 21, 33, 35, 37 and 59.)

[Mishra 2004] Arunesh Mishra, Minho Shin and WA Arbaush. *Context caching using neighbor graphs for fast handoffs in a wireless network*. In INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies, volume 1. IEEE, 2004. (Cited on pages 28, 36 and 44.)

[Montavont 2005] J. Montavont, N. Montavont and T. Noel. *Enhanced Schemes for L2 Handover in IEEE 802.11 Networks and Their Evaluations.* In IEEE 16th International Symposium on Personal, Indoor and Mobile Radio Communications, 2005. PIMRC 2005., volume 3, pages 1429–1434. IEEE, 2005. (Cited on pages 29 and 33.)

[Montavont 2006] J. Montavont and T. Noel. *IEEE 802.11 Handovers Assisted by GPS Information.* In Wireless and Mobile Computing, Networking and Communications, 2006. (WiMob'2006). IEEE International Conference on, pages 166–172, June 2006. (Cited on page 35.)

[Mota 2013] Vinicius FS Mota, Daniel F Macedo, Yacine Ghamri-Doudane and José Marcos S Nogueira. *On the feasibility of WiFi offloading in urban areas: The Paris case study.* In Wireless Days (WD), 2013 IFIP, pages 1–6. IEEE, 2013. (Cited on pages v, 18, 19, 20 and 45.)

[Mouton 2013a] Maximilien Mouton, German Castignani, Raphael Frank, Lara Codeca and Thomas Engel. *On the Evaluation of Make-Before-Break Handovers in Urban WiFi Networks for Moving Vehicles.* In Proceedings of IEEE WONS 2013 10th International Conference on Wireless On-demand Network Systems and Services, 2013, 2013. (Cited on page 4.)

[Mouton 2013b] Maximilien Mouton, German Castignani, Raphaël Frank and Thomas Engel. *MROAD: Scanning-free model-based roaming for 802.11 networks under vehicular mobility.* In Wireless Days (WD), 2013 IFIP, pages 1–8. IEEE, 2013. (Cited on page 4.)

[Mouton 2015] M Mouton, G Castignani, R Frank and T Engel. *Enabling vehicular mobility in city-wide IEEE 802.11 networks through predictive handovers.* Vehicular Communications, vol. 2, no. 2, pages 59–69, 2015. (Cited on page 4.)

[Nachtigall 2008] Jens Nachtigall, Anatolij Zubow and J-P Redlich. *The impact of adjacent channel interference in multi-radio systems using ieee 802.11.* In Wireless Communications and Mobile Computing Conference, 2008. IWCMC'08. International, pages 874–881. IEEE, 2008. (Cited on page 78.)

[Nadeem 2004] Tamer Nadeem, Sasan Dashtinezhad, Chunyuan Liao and Liviu Iftode. *TrafficView: traffic data dissemination using car-to-car communication.* ACM SIGMOBILE Mobile Computing and Communications Review, vol. 8, no. 3, pages 6–19, 2004. (Cited on page 9.)

[Nandan 2005] Alok Nandan, Shirshanka Das, Giovanni Pau, Mario Gerla and MY Sanadidi. *Co-operative downloading in vehicular ad-hoc wireless networks.*

In Wireless On-demand Network Systems and Services, 2005. WONS 2005. Second Annual Conference on, pages 32–41. IEEE, 2005. (Cited on page 9.)

[Narten 2007] Thomas Narten, Susan Thomson and Tatuya Jinmei. *IPv6 stateless address autoconfiguration*. 2007. (Cited on page 26.)

[Networks] Scalable Networks. *Qualnet Network Simulator, [online] Available*. http://www.scalable-networks.com. (Cited on page 55.)

[Niculescu 2007] Dragoş Niculescu. *Interference map for 802.11 networks*. In Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, pages 339–350. ACM, 2007. (Cited on page 21.)

[Ott 2004] Jörg Ott and Dirk Kutscher. *Drive-thru Internet: IEEE 802.11 b for" automobile" users*. In INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies, volume 1. IEEE, 2004. (Cited on pages 13, 18 and 22.)

[Pack 2005] Sangheon Pack, Hakyung Jung, Taekyoung Kwon and Yanghee Choi. *Snc: a selective neighbor caching scheme for fast handoff in ieee 802.11 wireless networks*. ACM SIGMOBILE Mobile Computing and Communications Review, vol. 9, no. 4, pages 39–49, 2005. (Cited on page 36.)

[Park 2004] Sang-Hee Park, Hye-Soo Kim, Chun-Su Park, Jae-Won Kim and Sung-Jea Ko. *Selective channel scanning for fast handoff in wireless LAN using neighbor graph*. In Personal Wireless Communications, pages 194–203. Springer, 2004. (Cited on pages 28 and 44.)

[Ramachandran 2006] Kishore Ramachandran, Sampath Rangarajan and John C Lin. *Make-before-break MAC layer handoff in 802.11 wireless networks*. In Communications, 2006. ICC'06. IEEE International Conference on, volume 10, pages 4818–4823. IEEE, 2006. (Cited on pages 3, 32 and 46.)

[Ramani 2005] Ishwar Ramani and Stefan Savage. *SyncScan: practical fast handoff for 802.11 infrastructure networks*. In INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE, volume 1, pages 675–684. IEEE, 2005. (Cited on page 30.)

[Robinson 2005] Joshua Robinson, Konstantina Papagiannaki, Christophe Diot, Xingang Guo and Lakshman Krishnamurthy. *Experimenting with a multi-radio mesh networking testbed*. In Proc. of WinMee, volume 5, 2005. (Cited on page 78.)

[Sadiq 2012] Ali Safa Sadiq, Kamalrulnizam Abu Bakar, Kayhan Zrar Ghafoor, Jaime Lloret and SeyedAli Mirjalili. *A smart handover prediction system based on curve fitting model for fast mobile ipv6 in wireless networks.* International Journal of Communication Systems, 2012. (Cited on page 33.)

[Shin 2004] Sangho Shin, Andrea G Forte, Anshuman Singh Rawat and Henning Schulzrinne. *Reducing MAC layer handoff latency in IEEE 802.11 wireless LANs.* In Proceedings of the second international workshop on Mobility management & wireless access protocols, pages 19–26. ACM, 2004. (Cited on page 29.)

[Soroush 2012] Hamed Soroush, Nilanjan Banerjee, Mark Corner, Brian Levine and Brian Lynn. *A retrospective look at the UMass DOME mobile testbed.* ACM SIGMOBILE Mobile Computing and Communications Review, vol. 15, no. 4, pages 2–15, 2012. (Cited on page 24.)

[Strobel 2013] Otto Strobel and IGI Global. Communication in transportation systems. Information Science Reference, 2013. (Cited on page 13.)

[Sun 2014] Weiping Sun, Okhwan Lee, Yeonchul Shin, Seongwon Kim, Changmok Yang, Hyoil Kim and Sunghyun Choi. *Wi-Fi could be much more.* Communications Magazine, IEEE, vol. 52, no. 11, pages 22–29, 2014. (Cited on page 13.)

[Taleb 2010] Tarik Taleb, Abderrahim Benslimane and Khaled Ben Letaief. *Toward an effective risk-conscious and collaborative vehicular collision avoidance system.* Vehicular Technology, IEEE Transactions on, vol. 59, no. 3, pages 1474–1486, 2010. (Cited on page 9.)

[Teng 2009] Jin Teng, Changqing Xu, Weijia Jia and Dong Xuan. *D-scan: Enabling fast and smooth handoffs in ap-dense 802.11 wireless networks.* In INFOCOM 2009, IEEE, pages 2616–2620. IEEE, 2009. (Cited on pages 28 and 44.)

[Velayos 2004] Héctor Velayos and Gunnar Karlsson. *Techniques to reduce the IEEE 802.11 b handoff time.* In Communications, 2004 IEEE International Conference on, volume 7, pages 3844–3848. IEEE, 2004. (Cited on pages 28, 31 and 44.)

[Villegas 2007] Eduard Garcia Villegas, Elena Lopez-Aguilera, Rafael Vidal and Josep Paradells. *Effect of adjacent-channel interference in IEEE 802.11 WLANs.* In Cognitive Radio Oriented Wireless Networks and Communications, 2007. CrownCom 2007. 2nd International Conference on, pages 118–125. IEEE, 2007. (Cited on page 21.)

[Wanalertlak 2011] Weetit Wanalertlak, Ben Lee, Chansu Yu, Myungchul Kim, Seung-Min Park and Won-Tae Kim. *Behavior-based mobility prediction for seamless* *handoffs in mobile wireless networks*. Wireless networks, vol. 17, no. 3, pages 645–658, 2011. (Cited on page 104.)

[Watkins 2012] Watkins. *Strategy Analytics: A Quarter of Households Worldwide Now Have Wireless Home Networks, [online] Available.* http://www.businesswire.com/news/home/20120404006331/en/Strategy-Analytics-Quarter-Households-Worldwide-Wireless-Home, 2012. (Cited on pages 11 and 15.)

[Zhang 2008] Jiansong Zhang, Kun Tan, Jun Zhao, Haitao Wu and Yongguang Zhang. *A practical SNR-guided rate adaptation*. In INFOCOM 2008. The 27th Conference on Computer Communications. IEEE. IEEE, 2008. (Cited on page 21.)

[Zhang 2014] Daqiang Zhang, Min Chen, Mohsen Guizani and Haoyi Xiong. *Mobility prediction in telecom cloud using mobile calls*. Wireless Communications, IEEE, vol. 21, no. 1, pages 26–32, 2014. (Cited on page 104.)

Acronyms

AC	Access Controller
AES	Advanced Encryption Standard
AP	Access Point
ARP	Address Resolution Protocol
ARPANET	Advanced Research Projects Agency NETwork
ASM	AP Selection Module
BATMAN	Better Approach To Mobile Ad-hoc Networking
BBM	Break-Before-Make
BSS	Basic Service Set
BSSID	Basic Service Set ID
CAN	Controller Area Network
CCH	Control Channel
CCK	Complementary Code Keying
CDB	Context DataBase
CDF	Cumulative Distribution Function
CDMA	Code Division Multiple Access
CM	Connection Module
CN	Community Network
CNo	Correspondent Node
CTP	Cabernet Transport Protocol
DAD	Duplicate Address Detection
DARPA	Defense Advanced Research Projects Agency
DCF	Distributed Coordination Function
DDM	Direction Detection Module

DHCP	Dynamic Host Configuration Protocol
DIFS	DCF Interframe Space
DMM	Direction Detection Module
DNS	Domain Name System
DSSS	Direct-Sequence Spread Spectrum
EDGE	Enhanced Data rates for GSM Evolution
ESSID	Extended Service Set ID
EWMA	Exponential Weighted Moving Average
FBSST	Fast Basic Service Set Transition
FEC	Forward Error Correction
FTD	Fuzzy-based Turning Detector
GIR	MIH Get Information Request
GPRS	General Packet Radio Service
GSM	Global System for Mobile communications
HA	Home Agent
HIP	Host Identity Protocol
HSPA	High Speed Packet Access
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IAPP	Inter-Access Point Protocol
ICMP	Internet Control Message Protocol
IP	Internet Protocol
IPv6	Internet Protocol version 6
ISP	Internet Service Provider
ITS	Intelligent Transportation System

L2	Layer-2
L3	Layer-3
LAN	Local Access Network
LGD	MIH Link Going Down
LoS	Line of Sight
LSE	Least Squares Estimator
LTE	Long Term Evolution
LTE-A	Long Term Evolution Advanced
LWAPP	Lightweight Access Point Protocol
MAC	Medium Access Control
MAG	Mobility Anchor Gateway
MANET	Mobile Ad-Hoc NETwork
MaxCT	Maximum Channel Time
MBB	Make-Before-Break
MCS	Mobile CrowdSourcing
MIH	Media-Independent Handover
MICS	Media Independent Command Service
MIES	Media Independent Event Service
MIIS	Media Independent Information Service
MIHF	Media-Independent Handover Function
MIMO	Multiple-Input and Multiple-Output
MinCT	Minimum Channel Time
MIP	Mobile IP
MIPv6	Mobile IPv6
MME	Mobility Management Entity

MPTCP	MultiPath TCP
MS	Mobile Station
NEMO	Network Mobility
ND	Neighbor Discovery
NG	Neighbor Graph
NLoS	Non Line-of-Sight
OBU	On-Board Unit
OFDM	Orthogonal Frequency-Division Multiplexing
P2P	Peer to Peer
PDR	Packet Delivery Ratio
PER	Packet Error Rate
PMIPv6	Proxy Mobile IPv6
PRNET	Packet Radio NETwork
PSK	Phase-Shift Keying
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
RP	Road Portion
RS	Road Segment
RSRP	Reference Symbols Received Power
RSS	Received Signal Strength
RSU	Road Side Unit
RSVP	Resource reSerVation Protocol
SCH	Service Channel
SDR	Software-Defined Radio

SI	Single Interface
SDM	Spatial Division Multiplexing
SLAAC	Stateless Address Autoconfiguration
SNR	Signal-to-Noise Ratio
SSID	Service Set ID
TCP	Transmission Control Protocol
TDMA	Time-Division Multiple-Access
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunications System
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
VANET	Vehicle Ad-hoc NETwork
VAP	Virtual Access Point
WAVE	Wireless Access in Vehicular Environments
WLAN	Wireless Local Area Network
xDSL	Digital Subscriber Line

Optimization of handover procedure between IEEE 802.11 access points under vehicular mobility in an urban environment.

Abstract: This thesis presents investigations performed on the IEEE 802.11 handover process in order to optimize the exploitation of exisiting IEEE 802.11 networks in vehicular communications. As IEEE 802.11 networks operate over a short range, a mobile station embedded in a vehicle faces frequent handovers leading to long disconnections. This work focused on two objectives. The first was reducing the IEEE 802.11 handover impact such that mobile users are allowed to be fully connected while they stay within the area covered by IEEE 802.11 networks. The second was optimizing the AP selection process in order to comply with the always best connected paradigm by identifying the best access points and the best handover location. These researches lead to the development of a context-based predictive handover mechanism that considers vehicle's trajectory, road topology, and network deployment information to decide the best handover location and candidate access points.

Keywords: IEEE 802.11, vehicular communications, handover optimization, predictive handover, context-awareness, simulation, experimental evaluation