Ph.D. Thesis

Self-Optimization of Radio Resources on IEEE 802.11 Networks

Ph.D. Candidate: Eduard Garcia Villegas
Ph.D. Advisor: Josep Paradells Aspas

Wireless Networks Group (WNG) — *Telematics Department*

Universitat Politècnica de Catalunya (UPC)

July 2009

Acknowledgements

At first glance, anyone may easily infer that the most evident results of the last five years of my life are an incipient alopecia and a loss of shape. I hope that the reader of this thesis would pay regard to other commendable achievements. These achievements are not the fruit of my effort alone since many people have helped me throughout these years. I would like to express my gratitude to those who stood beside me and supported me. Firstly I am particularly indebted to Professor Josep Paradells, who envisioned this thesis five years ago. He has guided me ever since to forge what you hold in your hands, and I wish to enjoy his verve for many years to come. I also like to acknowledge the advice and support of Rafael Vidal, which goes back to even before I "hatched" as an engineer. He has generously read and commented on almost every line I wrote during this thesis. I am grateful to Elena López for her "tech support" with *the simulator* - which I suggest to rename as Elena's WLAN incredible simulator (Elwis) -. Many thanks are to Marilet de Andrade for putting up with my "not-always-enjoyable" mood and for sitting with me in the same office without committing murder. I would also like to thank the reviewers (Jeff, Héctor Velayos and Pedro Ruiz) for their valuable and insightful comments that indubitably helped to improve the quality of this document.

There are many others that I hope they accept this *multicast* acknowlegement: to the members of the *WNG*, to my pupils (TFC/PFC), to the inhabitants of the lab 322, to the marauders of the C3, to my colleagues at the EPSC, to all the *guardiola* keepers,... THANK YOU!

Contents

List of Figures

List of Tables

Chapter 1

Introduction

It seems that the current step in evolution involves an attempt to spread a new primal need to those *Hominidae* bipedal primates known as *homo sapiens*. In other words, in the present era, people are inclined to be permanently “connected”. Since the dawn of the Information Age, engineers have been actively promoting this change. If the teleportation device still has not yet been invented while you are reading this, you will agree that one of the most recent developments that have pushed things along is the ability to access the Network of Networks (a.k.a. the Internet) wirelessly. Was it necessary? Certainly not; but it has opened up a new world of possibilities, since users now have the opportunity to access their home/office network resources without physically being present. A wireless access provides users with mobility, eases deployment where wired connections are not feasible, and reduces costs.

Early 3G deployments did not meet the expectations of a wireless broadband access, and wireless local area networks (WLANs) took advantage of this fact. WLAN technology evolved notably since the ALOHAnet (early ’70s) and has continued to do so right up to the present. Primitive development included industry-specific solutions and proprietary protocols, but at the end of the ’90s these were replaced by standards, which constituted a milestone in the explosion of WLANs.

WLANs, primarily the various versions of IEEE 802.11 standards, and more precisely, those operating in infrastructure mode using the distributed coordination function (DCF), are nowadays the most popular technologies for providing broadband radio access to IP networks, whether to extend Small Office/Home Office (SOHO) network LANs or to provide Internet access in public places (such as airports, hotels and even coffeehouses). Moreover, with the appearance of “Wi-Fi Certified”[1] products, different competitive brands of WLAN devices are interoperable at a basic level of service. Therefore, unlike cellular phones, any standard Wi-Fi device will work almost anywhere in the world.

Nevertheless, some say that with the advent of 4G (LTE, WiMAX) the end of the IEEE WLANs is approaching, but according to AT&T reports [38], the number of Wi-Fi

[1] http://www.wi-fi.org

connections in the first quarter of 2009 is three times greater than in the first quarter of 2008. The increasing saturation of laptop-style computers (with built-in Wi-Fi interfaces), netbooks, smartphones (the popularity of the iPhone, as well as Wi-Fi enabled Blackberrys), are driving Wi-Fi use nowadays.

However, the increasing density of WLAN access points has started to reveal the negative effects and shortcomings of the original IEEE 802.11 standards. One of its key success factors, the use of unlicensed Industrial Scientific Medical (ISM) frequency bands, is at the same time one of its major drawbacks. "Unlicensed" means that these frequencies are freely available to the general public. On the other hand, such frequencies are defined within a small portion of the spectrum and are usually shared among several users, as well as being used by many different technologies (e.g. cordless phones, wireless cameras, Bluetooth, ZigBee, etc.). Besides, the type of medium access defined by the IEEE 802.11 (CSMA, or "listen before talking") requires a special attention to all interference issues, as detailed throughout this document.

In this scenario, IEEE WLANs are unable to exploit all their potential. However, intelligent radio resource management (RRM) policies could be applied to minimize the harmful effects of interference and an uneven load distribution. Due to the inherent flexibility of WLANs, the use of unlicensed frequency bands and their (in general) non-commercial nature, network administrators and manufacturers avoided the complexity involved in other cellular technologies (e.g. GSM, UMTS, etc.) and therefore have not paid much attention to radio resource management. In contrast, the study and development of efficient RRM strategies has been the concern of many researchers, who have demonstrated that even the simplest mechanisms are able to improve notably the performance of a legacy IEEE WLAN. In this regard, the works by Arunesh Mishra, Héctor Velayos [284], Gunnar Karlsson, Petri Mähönen, Janne Riihijärvi, Marina Petrova, Andreas Eisenblätter and Iana Siomina [269] deserve a special mention.

This thesis explores ways of understanding the performance issues that are endemic to IEEE 802.11 WLANs, as well as ways of minimizing these negative effects by means of radio resource management. In fact, these problems are not new and have been studied extensively since the advent of mobile communications networks, but the particular characteristics of the 802.11 WLANs, as discussed in depth throughout this document, require new approaches.

1.1 Radio Resource Management (RRM)

Planning is the traditional tool used in radio networks for tackling the problems of avoiding interference, providing enough coverage despite the propagation phenomena, and guaranteeing a minimum quality, or at least, a minimum capacity. Planning consisted in a deep study of the environment, including orography, the presence of obstacles, traffic demand points, measured interference, etc. Then, an offline calculation provides the configuration

parameters for optimizing the network resources under those circumstances. However, as discussed earlier, the dynamism of the WLAN environment is such that the initial settings provided during the planning phase will need frequent revision.
Given the *dynamic* nature of radio propagation, and due to user mobility, the parameters that define the wireless environment vary in time. Therefore, in order to maintain the required performance, we should find mechanisms that allow a dynamic reconfiguration of the network in response to the changes in the environment. In large-scale deployments, or in a rapidly varying context, a manual configuration makes no sense, unless one has a horde of trained Oompa-Loompas at one's disposal. Since this is not usually the case, the best way of obtaining a constant adaptation lies in providing the network elements with the ability to detect changes in the environment, and change their settings *automatically.*

With the ongoing proliferation of wireless users and access points, it often happens that different wireless networks coexist in the same area. These networks may belong to different administrators, whose only relation is merely the use of a common medium: the unlicensed radio spectrum. This situation hinders the implementation of *centralized* RRM mechanisms. A centralized approach is based on a central entity that is aware of the whole network and can remotely control any configuration parameter of any network element. However, no central manager will be allowed to handle these devices if such elements are owned by separate entities. In this case, a more suitable approach involves the different elements having enough *intelligence* to take decisions by themselves; that is, RRM management would be carried out in a *distributed* manner. Even so, the centralized approach remains the most widely accepted management paradigm. Depending on the scenario, centralization could offer better scalability, and would be preferable due to its higher simplicity. In consequence, it still deserves our study and attentions

A set of elements that are able to take decisions by themselves, solely by taking into account the portion of the "world" they see, can aspire only to obtaining local optimizations [247]. In contrast, as we all learned from [170], *cooperation* among the different elements in conflict contributes to a global optimization, which improves the resources available for every individual. In other words, in order to obtain global improvements that benefit all the elements in competition, these elements should work collaboratively. The cooperation among network devices entails new issues: where does this collaboration take place? And at what level? How are these elements coordinated, and how can they communicate? Recall that these network elements may belong to different domains.

To sum up, we build our scenario of interest with *intelligent* Wi-Fi devices capable of *cooperating* either in a *centralized* or a *distributed* manner, in order to make a better use of the shared and scarce *radio resources.* There still remains a further keyword: *open source.* An efficient RRM mechanism may be the panacea from a theoretical perspective, but it is useless if it cannot be put in practice. Our aim is to develop RRM mechanisms for real IEEE 802.11 WLANs, and therefore we need an "open" access to the bowels of these devices. This led us to focus our development on Wi-Fi equipment that works on open source platforms, i.e. hardware platforms running open source Operating System (OS), with open source drivers (more details in Section 1.2).

1.1.1 Dynamic RRM mechanisms for IEEE 802.11 WLANs

RRM mechanisms in the field of IEEE 802.11 WLANs are basically intended to reduce contention and interference. This reduction is translated into an improved Quality of Experience (QoE), as perceived by the users. To this end, RRM should provide efficient channel allocation mechanisms, modulation selection algorithms, power control and load balancing. Finding optimal AP placements is an issue that is often covered in the literature. This optimization is combined with frequency and transmitted power assignments, and has a threefold aim: reducing the number of APs required to cover a given area, balancing load in congested areas, and reducing inter-cell interference. However, AP placement decisions are mostly taken during the planning phase, with no chances of being re-visited once the network is deployed, even though the changing environment might require it. For these reasons, we will not consider AP placement studies within the framework of our dynamic RRM strategies[2].

1.1.1.1 Rate selection

The IEEE 802.11 standards define different sets of modulations for the Physical Layer (PHY). For example, in IEEE 802.11a there are four modulations (BPSK, QPSK, 16-QAM and 64 QAM), and one convolutional coder providing different coding rates for bit error correction. By combining modulations and coding rates, IEEE 802.11a offers eight different modes with different physical rates, from 6 to 54 Mbps (see Table 2.2). Logically, a higher physical rate will provide a better application throughput. However, this is not always true; a faster modulation scheme is also more sensitive to noise and interference, and will consequently require a higher Signal to Interference and Noise Ratio (SINR). In other words, fast modulations will provide a good net throughput when the SINR is high, but their performance will degrade as the SINR is decreased. There are some points where slow modulations compensate for their lower physical rates by an increased robustness against noise and interference, thereby improving the frame loss rates obtained by their "fast" counterparts. In fact, there is a range of SINR values for which a given modulation outperforms the others (details in Sections 3.2 and 4.3.1).

An automatic rate selection mechanism (a.k.a link adaptation) allows a Wi-Fi transmitter to select the most appropriate mode according to the status of the channel. Ideally, these mechanisms should adapt the PHY rate in transmission depending on the signal quality measured at the receiver. Obviously, this information is known only by the receiver. Will the transmitter be able to assess the receiver's reception quality? Should

[2] Even though it may be feasible, it is not easy to integrate AP placement within a dynamic RRM strategy, since it would possibly require moving the APs on demand (e.g. see http://www.wifibot.com)

the receiver share this information with the transmitter? The new IEEE 802.11h features will assist in these decisions. This discussion is further extended in Section 4.3.3, where a model for a simple algorithm is also provided.

1.1.1.2 Load balancing

As explained later in Section 2.1, WLANs are built according to two architectures: ad hoc and infrastructure-based. The second is the most common nowadays. It relies on the use of wireless access points (APs) - Wi-Fi devices with special functions - which are connected to a wired infrastructure. Client stations access the network through these APs. An AP establishes an 802.11 cell and coordinates all communications (air-to-air, air-to-wire and wire-to-air) that take place in that cell. The boundaries of the cell are basically defined by the AP's transmitted power and sensitivity, and by propagation effects.

An area is covered by several APs, and the client stations chose one or another according to received power measurements. Since bipedal primates, to whom most WLANs are devoted, are gregarious (and tend to concentrate in time and space), a few APs will usually have to serve many users, while other distant APs will remain under-utilized. This situation can be compounded by the presence of fixed devices, such as Wi-Fi storahe disks, network printers, etc. If many users are using the same AP, their share of the common medium will be small, thus getting a poor service. On the other hand, users connected through an under-utilized AP will get more resources, which is an advantage for the unsociable.

Load balancing techniques in WLANs try to equalize the utilization of all the APs, so that the service provided to the users is shared fairly. These techniques need first to define the concept of load and fairness, and then take the appropriate actions. These actions can be of different nature. One first approach consists in choosing an AP according to load-based criteria, and not only power measurements. Note that choosing a less loaded (and distant) AP could entail a worse signal quality and thus a slower modulation. Hence, rate adaptation has a notorious impact on load balancing and should be studied jointly. APs could also accept or deny new client associations depending on their current load. Finally, load balancing can be combined with AP's transmitted power control: reducing or increasing the cell dimensions the AP serves less or more clients. All these issues are discussed in depth in Chapter 6.

1.1.1.3 Frequency channel allocation

IEEE 802.11 WLANs use unlicensed frequency bands known as ISM bands. IEEE 802.11b and 802.11g standards define 11 to 14 channels within the 2.4 GHz ISM band, depending on the regulatory domain (c.f. Section 3.3.3). Channel 1 is centered at 2.412 GHz, channel 11 (highest available channel in North America) at 2.462 GHz, channel 13 (highest available channel in Europe) at 2.472 GHz, and channel 14 (available in Japan) is at

2.484 GHz. Except for channel 14, the center frequency distance between consecutive channels is 5 MHz.

Besides the center frequencies, the IEEE 802.11 specifies a spectral mask defining the permitted distribution of power across each channel. The signal must be attenuated by 30 dB or more from its peak energy at ±11 MHz from the center frequency. Therefore, channels are effectively 22 MHz wide. The first consequence is that only three or four channels do not overlap, typically 1, 6 and 11 (in the Americas), or 1, 5, 9 and 13 (in Europe). In other words, the 2.4 GHz ISM band provides room for only three or four nearby simultaneous transmissions without interference.

In infrastructure-based WLANs, a channel is assigned to all the elements associated to a given AP, so that the interference between neighboring cells is minimized. More precisely, a channel is assigned to the AP, and the client devices will use the channel configured for the AP. Usually, this channel is pre-configured or selected randomly at start up. The most sophisticated of today's APs scans the medium once (before its cell is set), or periodically, seeking for activity on the different channels; then the least congested channel (or the channel containing less APs) is chosen. This approach has several flaws; for example, it does not take into account the interference seen by the client devices, which now have the ability to send interference reports through IEEE 802.11k messages. More sophisticated strategies are discussed in Chapter5.

Many channel allocation strategies are intended to assign different orthogonal channels to neighboring cells. However, in dense scenarios this problem is not solvable, since three channels are often not enough, and two or more cells will have to contend for the same portion of the spectrum. This is known as *co-channel interference*: interfering transmissions produced in the same channel. Co-channel interference may prevent a given station from accessing the channel, due to the CSMA access used in 802.11 WLANs. On the other hand, an interfering transmission on a partially-overlapping channel (or adjacent channel), is called *adjacent-channel interference*. Depending on its intensity and on the channel distance, this interference could be added to the noise floor, with the resulting increase in the number of reception errors. Whether or not co-channel interference is preferable to adjacent-channel interference depends on many parameters. The effects of both types of interference need to be studied before deciding on a channel allocation strategy. This study is carried out in Sections 3.3 and 4.3.

IEEE 802.11a WLANs use the 5 GHz ISM band, where up to 19 non-overlapping channels are available. In addition, this band suffers less radio interference, as compared to the 2.4 GHz band. In consequence, channel allocation in 802.11a networks is not as critical as in 802.11b or 802.11g. However, the presence of RADAR systems in the 5 GHz ISM, motivated the definition of IEEE 802.11h, which adds dynamic frequency selection capabilities to Wi-Fi devices. Furthermore, the emergence of 802.11n devices, which can improve their capacity by making use of two channels simultaneously, is increasing the interest in frequency assignments in the 5 GHz band.

1.1.1.4 Power control

Even though IEEE WLANs use “free” frequency bands, they should conform to radio regulations and limit their transmitted power in order to limit the interference caused to other systems or networks. The recent IEEE 802.11h adds features that ease the adaptation of transmitted powers to the particular limits specified for a given region.

In dense scenarios, IEEE 802.11 WLANs are interference-limited due to the scarce frequency spectrum, and therefore a reduction in transmitted power is useful in minimizing such interference. However, reducing the transmitted power entails a SINR degradation in reception, which is in turn translated into an increased frame error ratio. Moreover, if automatic rate selection comes onto the scene, a slower modulation will be chosen, and the net throughput of the cell will fall.
On the other hand, reducing transmitted power also means reducing power consumption, which is a key issue when client stations come in the form of portable devices powered by batteries (e.g. laptops, handhelds, etc.).

Furthermore, as mentioned earlier, power control and load balancing can be studied jointly. Over-loaded APs could reduce their cell dimensions (i.e. reduce transmitted power), while under-utilized neighboring APs increase their radius (i.e. increase transmitted power) in order to attract distant clients and cover possible coverage gaps. This technique is known as cell breathing, and is explained in Section 6.2.3.

1.2 Implementation platforms

The research developed within this thesis has been deliberately implementation-oriented. As mentioned, the objective is not limited to the investigation of algorithms, architectures, paradigms, etc. With the firm conviction that these mechanisms have real application and can be put into practice, part of this thesis is devoted to working with real off-the-shelf devices, paying special attention to the 802.11 APs.

Although the first practical implementations were developed on Linux-based PCs, those developments were later ported to different commercial AP platforms: *Linksys WRTG54G*, *4GSysteme AccessCube* and *Futurlink Wilico B200*. It is well known that all these platforms support a Linux embedded OS, and thus their firmware can be “easily” replaced and upgraded with custom-made applications.

The story of embedded open-source 802.11 APs starts with the Linksys WRT54G. This wireless router, released in December 2002, was one of the first commercial products that adhered to the IEEE 802.11g amendment. In that same year, some developers found that its firmware was based on Linux components[3]. Because Linux is released under the GNU General Public License (GPL), the terms of the license obliged Linksys to make available

[3] http://lkml.org/lkml/2003/6/7/164 the original post that appeared in the Linux Kernel Mail Archive

the source code present in the WRT54G firmware. This fact allowed developers to learn how the hardware was accessed, and motivated many Linux enthusiasts to develop new add-ons or even new firmware distributions made from scratch.

Embedded computers, such as commercial WLAN APs, usually have extremely limited resources (memory, processing capacity, etc.). These computers are thus not powerful enough to run a compiler or a development environment. Since debugging and testing may also require more resources than those available on an embedded system, these processes actually take place on another computer (e.g. a PC). This technique is known as cross-compilation.
With cross-compilation, *build environment* and *target environment* are located in different devices. The first one is where the development and debugging is carried out. Furthermore, this system is also responsible for producing the binaries that will be run on the *target environment*, that is, the resource-limited embedded system. To this end, the *build environment* must be provided with a complete toolchain. The toolchain consists of a complete set of software utilities that include, among other things, a compiler system, and a collection of libraries capable of running on different computer architectures. In Linux environments, the compiler system is based on GNU Compiler Collection (GCC), mostly used to compile C and C++ code, although other languages are supported. Since all of our codes are programmed in *C*, the standard library used is the GNU C Library (glibc). However, glibc may not be suitable for embedded devices, since it requires many resources. Instead, μC Library (uClibc) is much smaller than glibc, and besides, it supports nearly all applications designed for glibc.

Our firsts attempts were developed and implemented in PC Card-enabled desktops or laptops running a Linux-based OS. Under Linux, a Wi-Fi device can be managed from user-space through the Jean Tourrilhes's Wireless Extensions (WEXT) [277]. The WEXT provide an Application Programming Interface (API) that allows manipulation of any wireless networking device in a standard and uniform way, that is, independently of the device or the driver. It defines two ways of accessing the device: via the */proc* pseudo file system, and via Input/output control (ioctl) calls. This easy-to-use interface enables access to firmware/driver level statistics to user-space applications. However, not all drivers exploit the potential of WEXT, thus reducing the set of interesting devices to those having a HostAP [217] or MadWiFi [198] enabled hardware. For example, in [64], channel and transmitted power configurations were tested. In [131] and [146] we studied the statistics made available by the driver.
However, wireless drivers developers are moving to the new *mac80211* stack[4], included in kernels from 2.6.22, although it is still experimental and a lot of work remains to be done. For example, at its current development stage, the new stack does not yet support channel or transmitted power settings on a per-packet basis. This is still done through old netlink and ioctl calls.

The Wilico B200 platform (a.k.a. Infopoint) is based on an x86 architecture, and

[4] http://wireless.kernel.org

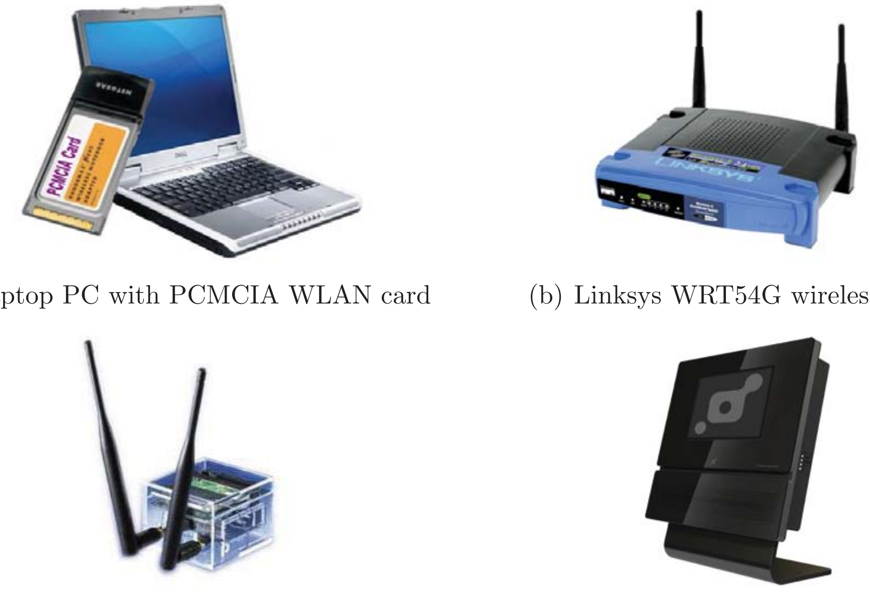

(a) Laptop PC with PCMCIA WLAN card

(b) Linksys WRT54G wireless router

(c) 4G Systems AccessCube

(d) Futurlink Wilico B200A

Figure 1.1: Linux-based APs used in our implementations

therefore the cross-compilation process is not required, even though development and debugging actually take place in a desktop PC.

The WRT54G and variants (WRT54GS, WRT54GL, etc.) are IEEE 802.11b/g devices from Linksys (see Fig. 1.1(b)). These family of wireless APs are typically used as routers to provide wireless and wired access to Internet in a SOHO environment. It provides a four Ethernet port switch and a WAN (DSL) interface. Newer models (e.g. WRT54G3G) also include a PCMCIA slot for use with a cellular based Network Interface Card (NIC). The development of applications for this platform was studied in [263]. This work also analyzes the wide variety of available firmwares, concluding that *OpenWRT*[5] offers the best alternative.

Similar to the WRT54G, the 4G Systeme AccessCube (a.k.a. meshCube) is based on a Microprocessor without Interlocked Pipeline Stages (MIPS) architecture, though it is provided with improved features (see Table 1.1). Moreover, this platform offers two ways of expanding the hardware: a Mini-PCI bus and an Inter-Integrated Circuit (I^2C) bus. The first is used to add one or more wireless NICs (e.g. WI-FI, Bluetooth, etc.). The I^2C bus allows the attachment of low-speed peripherals, such as sensors, a servo controller, etc. In [139], we studied the development of applications for this platform, based on the *OpenEmbedded*[6] framework.

[5] http://openwrt.org

[6] http://www.openembedded.org

	Laptop	**WRT54G**	**AccessCube**	**Infopoint**
CPU Arch. (MHz)	x86 900-1600	MIPS 125-240	MIPS 640	x86 600-1000
RAM (MB)	500-1000	8-32	64	256-500
Storage (MB)	SATA/ATAPI 40000-80000	Flash 2-8	Flash 32	C-Flash 4000-8000
WLAN chipset	Prism/Atheros	Broadcom	Prism/Atheros	Prism/Atheros
Firmware/OS	Debian	OpenWRT-EH	Nylon	Debian

Table 1.1: Features of different AP platforms

1.3 Scope and structure of this thesis

This thesis is composed of eight chapters, the first of which corresponds to the introduction. This chapter first depicts the "habitat" of wireless packet networks, more precisely, Wireless Local Area Networks based on the IEEE 802.11 set of standards. After discussing some of the radio resource sharing and performance issues raised by this kind of network, we briefly describe some existing solutions together with the different approaches provided in this thesis. Finally, this section outlines the remaining of the document, paying a special attention to our contributions.

Following this introductory chapter, Chapter 2 provides basic background on the IEEE 802.11 standards: its architecture and topologies, the services and operations defined for the Logical Link Control (LLC) and Medium Access Control (MAC) layers, and the main characteristics of the most spread physical layers. The operation of 802.11 technology is described at a level that is essential to the understanding of the contents of this thesis. Finally, we try to shed some light on the endless alphabet soup that the IEEE 802.11 Working Group (WG) seems determined to make more confusing than a David Lynch movie.

Chapter 3 deals with the WLAN environment in depth, in order to characterize the radio channel from an IEEE 802.11 transmitter or receiver's point of view. First of all, Chapter 3 reviews the propagation models used in WLAN scenarios. Next, error models for the most common WLAN physical layers (PHYs) are given. This chapter is also intended to provide a new analysis on the effects of interference, exposing one of the contributions of this thesis. Knowledge of all these phenomena is indispensable if we are seeking a better utilization of the radio resources, and it is therefore essential that the simulation tools used take them into account. For these reasons, a study on the different available simulation environments is provided.

Logically, all the phenomena studied in Chapter 3 affects the performance of a WLAN. Those harmful effects are mainly translated into a throughput degradation or into a loss of capacity. In consequence, throughput is a reliable performance metric that can be used to assess the radio resource management. Chapter 4 is then devoted to studying throughput and capacity models, paying special attention to their complexity and the way they capture the radio phenomena studied in Chapter 3. There are models for single user, or single cell scenarios, but in the literature none was found for modelling the capacity of a multiple cell network. In this regard, this chapter includes the second contribution of this thesis: a throughput model for large WLANs.

Chapter 5 is devoted to the frequency assignment problem in the WLAN context. The approach detailed in this chapter is based on the capacity model for multi-cell WLANs presented in Chapter 4, which in turn takes all the effects studied in Chapter 3 into account. Our contribution here is twofold: first, we present a scheme that is able to make use of all the available spectrum, instead of the three traditional non-overlapping channels; and second, the low operational complexity of the system makes it suitable for running in low-featured devices in a distributed manner. The improvements achieved by our approach are evaluated by means of simulations and measurements in a small testbed.

In Chapter 6, a discussion is settled around the concept of load. In the case of IEEE 802.11 WLANs, this concept of load is not as trivial as in other networks, as argued in the first section of Chapter 6. This discussion led us to our contribution: a new load metric which effectively captures all the phenomena affecting the capacity of a WLAN cell (e.g. number of users, offered traffic, interference, etc.). Once the concept of load is defined, we contribute with two load balancing solutions: first, a client-driven scheme; and second, a cell breathing algorithm.

In order to put our radio resource management mechanisms in practice, it is necessary to integrate these mechanisms into a management architecture that enables the exchange of information among the involved devices. In Chapter 7, we take the opportunity to review the newest trends on this matter, including the most interesting (in our opinion) and yet less known IEEE 802.11 amendments (802.11h, 11k and 11v), together with the newly released CAPWAP protocol, result of recent IETF efforts. This chapter is also intended to describe our contributions in this area: a centralized and a distributed radio resource management scheme for IEEE 802.11 WLANs.

The final chapter presents the concluding remarks. Chapter 8 is then devoted to summarizing all the achievements of this thesis, and also to present possible ways of extending this work in the near future.

1.3.1 Contributions and publications

To sum up, the original contributions, result of our work during the development of this thesis, are identified in the following six points:

- The study and characterization of interference in the particular field of IEEE 802.11 networks. The results of this study were published in [129].

- A capacity model for large WLAN deployments that takes both the effect of carried traffic and inter-cell interference into account. The model also includes the effect of rate adaptation. These capacity estimations are useful to evaluate the benefits of a radio resource management mechanism. The model was presented in [130] and [135]. The work in [210] is focused on the effects of transmission errors.

- A frequency assignment mechanism that makes use of all the available spectrum, since it takes the effects of both co-channel and adjacent-channel interference into account. This idea was first presented in [134], although it was in an embryonic form. In [133] we discuss some implementation issues, and finally in [138] full details are given.

- A new load metric, useful in the context of IEEE 802.11 WLANs. The load metric is based on an algorithm that we published in [131]. A deeper evaluation was provided later in [126].

- Two load balancing mechanisms. Both ideas were first discussed in [136]. The client-driven approach was later detailed in [128], and the cell-breathing mechanism was presented in [137].

- The design and development of two radio resource management architectures. A centralized architecture, as part of the UAMN project [243, 244], and a distributed scheme, published in [127].

Chapter 2

IEEE 802.11 Wireless LANs

IEEE 802.11 is a set of standards for computer communications in a Wireless LAN (WLAN), developed by the Institute of Electrical and Electronics Engineers (IEEE) LAN/MAN Standards Committee (IEEE 802). The scope of these standards is to define one MAC and several PHY specifications for wireless connectivity for fixed or moving stations (STAs). As with other IEEE 802 based protocols (e.g. 802.3 and 802.5), the primary service of the 802.11 standard is to deliver MAC Service Data Units (MSDUs) between peer LLC entities.

This chapter provides an overview of the functions specified by the different IEEE 802.11 standards, paying special attention to the most common processes and the most popular PHYs.

2.1 The Basics

This section introduces the basic elements that are present in an IEEE 802.11 network, the possible topologies, depending on the configuration and capabilities of those elements, and the basic services that the MAC layer implements to provide a WLAN that enables station mobility transparent to higher layers.

2.1.1 The Elements

The most basic element in a WLAN is the **Station (STA)**. Any device that contains an IEEE 802.11-conformant MAC and PHY interface to the wireless medium is considered a STA. The functions of the 802.11 standard are implemented as part of the hardware in a radio NIC, and as software pieces that run either in the OS kernel (interface drivers) or in user space (applications). A STA is called **QoS-STA** if it implements the QoS functions introduced in [22].

An **Access Point (AP)** is an entity that implements the functionalities of a STA and besides, as well as providing access to the distribution services through the wireless medium, for associated (or *client*) STAs. If an AP is present, all communications between two client STAs are done through the AP, even though both STAs were close enough so as to contact directly. If PCF access is used, the AP is also the point coordinator, in charge of performing the polling of PCF client STAs (see 2.3 for more details).

The basic building block of an IEEE 802.11 WLAN is known as Basic Service Set (BSS). A BSS is a set of connected STAs. Strictly speaking, membership in a BSS does not imply that wireless communication with all other members of the BSS is possible. However, a BSS is often used to describe an AP and its associated STAs; this architecture is also known as *infrastructure BSS*.
Each BSS is identified by a BSS Identification (BSSID). In an infrastructure BSS, the BSSID is the MAC address of the AP. Otherwise, the BSSID is a locally administered MAC address generated from a 46-bit random number.

The **Distribution System (DS)** is used to interconnect a set of BSSs, typically through a wired infrastructure (e.g. an Ethernet LAN). The DS or backbone can also be completely built wirelessly through the Wireless Distribution System (WDS).

2.1.2 The Topologies

Using the BSSs as the building blocks of an IEEE 802.11 LAN, two different topologies are supported, as depicted in fig. 2.1: Independent Basic Service Set (IBSS) and Extended Service Set (ESS).

2.1.2.1 IBSS

When a BSS forms a self-contained network without infrastructure and with no access to a DS available, it is called IBSS, also known as ad hoc network in the IEEE 802.11 standard. In an IBSS, any STA can establish a direct communication with any other STA belonging to the same BSS, without the requirement of using a centralized intermediary acting as an AP.
For proper operation of this mode, all STAs should be in reach of each other. If this is not the case, new layer 2 (future 802.11s amendment) or layer 3 functionalities (wireless ad hoc routing protocol [100]) should be implemented in order to enable communications between all pairs of STAs.

2.1.2.2 ESS

The ESS consists of multiple BSSs using a common distribution system; an ESS is perceived by the LLC sublayer as a single large BSS. The basic services provided by the

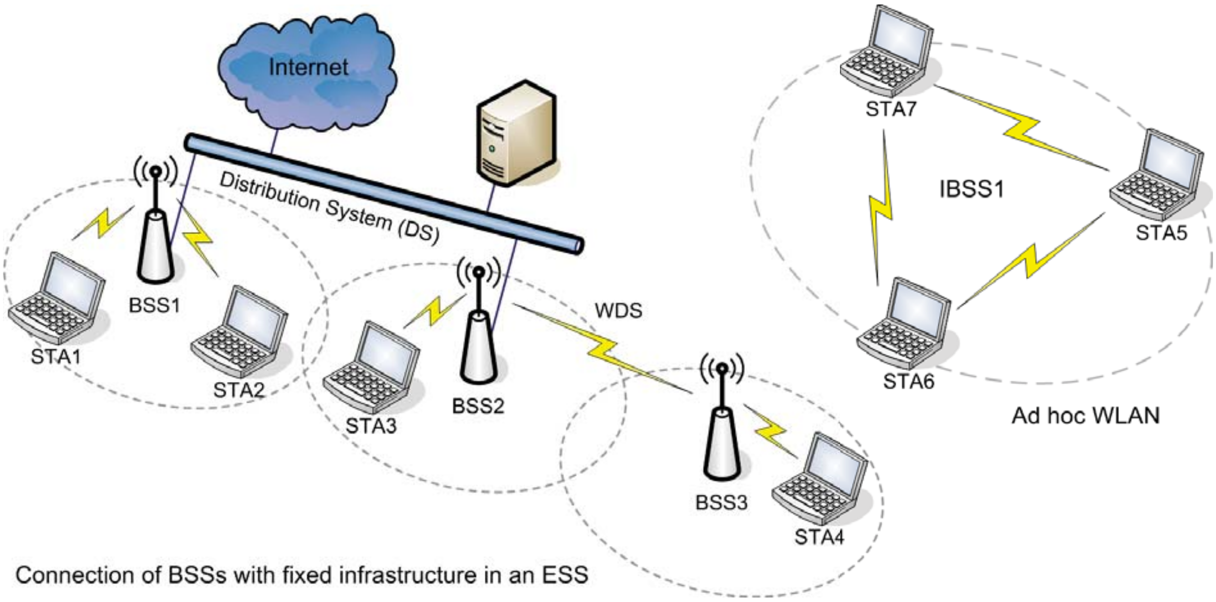

Figure 2.1: IEEE 802.11 WLAN supported topologies

common distribution system are intended to transport frames between stations that are not in direct communication with each other. That is to say, these services include transport between the APs within an ESS. An ESS also provides gateway access for wireless users into a fixed wired network, such as a Local Area Network (LAN) or the Internet. Furthermore, a client STA may move from one BSS to another (within the same ESS) transparently to the LLC. All APs in the same ESS are set according to a common ESS Identification (ESSID) which identifies the network. A WDS can also be used for hard-to-wire locations or for extending the BSS's coverage area.

2.2 The Services

The standard IEEE 802.11 defines the services that the DS and the STAs should implement. The complete list of MAC services is shown in Table 2.1. The table also indicates whether these services were part of the original definition or added in subsequent amendments. In this section we briefly describe these services, assuming an ESS topology, which is the preferred environment in this thesis.

2.2.1 Distribution of messages

Distribution of data messages between STAs is the primary service provided by an IEEE 802.11 network. It relies on the MSDU delivery function present in all STAs. MSDU delivery is detailed in section 2.3.

Service Name	**Provided by**	**Default**
Authentication	STA/DS	(11i)
Association	DS	yes
Deauthentication	STA/DS	(11i)
Disassociation	DS	yes
Distribution	DS	yes
Integration	DS	yes
Data confidentiality	STA	(11i)
Reassociation	DS	yes
MSDU delivery	STA	yes
DFS	STA	(11h)
TPC	STA	(11h)
Higher layer timer synchronization	STA	(11e)
QoS traffic scheduling	STA/DS	(11e)

Table 2.1: IEEE 802.11 architectural services

Refer to the ESS depicted in Fig. 2.1 and consider a message from STA1 to STA3. The distribution of the message is as follows: STA1 sends the message to its AP, the AP gives the message to the DS. It is the job of the DS to deliver this message to the appropriate destination, in this case, STA3's AP. Finally, the "output" AP accesses the wireless medium to send the message to the final destination (STA3). This operation may require the frame to be forwarded across different layer 2 technologies; the adaptation needed to deliver a frame to a non-802.11 network is carried out by the integration service, present in the APs. How the message is distributed within the DS is beyond the scope of the standard. However, IEEE 802.11 services are required to provide the DS with enough information to be able to forward the message towards its correct destination. This information is provided by the association related services (see 2.2.2).
In the previous example, a message is sent to a STA that belongs to a different BSS. If the message had been intended for STA2, the "input" and "output" AP would have been the same. Nevertheless, the distribution service would have been invoked anyway. Note that all data messages sent by any client STA in an infrastructure BSS must always be sent to the BSS's AP, even though source and destination STAs are in range of each other.

With WDS, the wireless link between APs uses the same radio as that used for communication with client STAs. Single radio APs are cheaper, but the cells of both APs as well as the wireless link all share the same channel, which reduces drastically the network capacity since some frames may be sent to the wireless medium three times: in Fig. 2.1 one frame transmitted by STA3 addressed to STA4 is first sent by STA3 to BSS3's AP; the AP uses the same channel to forward the packet to BSS4's AP, who will finally transmit it again to STA4. This inefficiency can be alleviated if APs have two (or more) radios using different channels.

The QoS traffic scheduling service is integrated within the distribution service of QoS-enabled networks to provide intra-BSS QoS frame transfers (see 2.5).

2.2.2 Association services

Before a STA is allowed to send data frames via an AP, it first has to become associated with the AP. This operation provides a STA ↔ AP mapping that is used by the DS to accomplish its distribution service as described in the previous section (2.2.1). For this reason, a client STA should be associated with no more than one AP, while an AP may be associated with many STAs at the same time.

In order to support BSS transition mobility, the DS also offers the reassociation service. This service is used to keep STA ↔ AP mappings updated as the STAs move from BSS to BSS within an ESS. Both association and reassociation services are always initiated by the client STA. The STA sends an *Association* (or *Reassociation*) *Request* management frame to the selected AP, and the AP responds with an *Association* (or *Reassociation*) *Response*. This client-driven behavior, detailed in 2.2.2.1 and 2.2.2.2, is the source of many administrators' headaches (i.e. load balancing) as discussed in this thesis.

Disassociation is used to terminate an existing association. As a consequence, the existing association information is removed. Unlike association and reassociation, a disassociation can be initiated either by clients or by APs, and cannot be refused since it is not a request but rather a notification.

2.2.2.1 Scanning

As explained before, when a client STA in infrastructure mode is powered up, it needs to establish an association with an AP. The scanning process allows the client STA to become aware of its environment, and hence the STA is able to chose the "best" AP in range. All APs are announced in an area by broadcasting *Beacon* frames periodically. These management frames contain information on the APs' capabilities and status that is useful for the potential clients. STAs usually chose the "best" AP based on Received Signal Strength Indicator (RSSI) measurements of *Beacon* frames. That is to say, client STAs will request an association with the AP whose *Beacons* are received with more energy during the scan process.

Scanning can be either active or passive. Stations that use passive scanning listen to each channel and wait for *Beacon* frames to identify all APs within range. Unlike in active scanning, STAs do not generate any frame during a passive scan. This works well for networks with few channels and short *beacon* intervals. Active scanning is a faster process, but requires a signaling exchange. To satisfy regulatory requirements (active scanning is prohibited in some frequency bands and regulatory domains), it may be required that an initial scan is always passive.

The active scan procedure is as follows. For each supported channel, the station willing to associate with a new AP broadcasts a *Probe Request* frame. Any AP shall respond with a *Probe Response* to the address of the STA that generated the *Probe Request*, whenever the requested ESSID matches the AP's configured ESSID.

2.2.2.2 Roaming

The DS has to support user inter-BSSS mobility within the ESS. When the client STA perceives that the quality of the communication with its current AP degrades due to either its mobility or to the presence of interference, it tries to find a better AP candidate through a new scan. This is called roaming, which goes as follows [246].

A station keeps track of the *Beacon* frames received from its current AP. When the quality of such *Beacons* drops below the cell search threshold ($10 < CS_{Th} < 30$ dB), the STA initiates an active scan (if possible) and sends out *Probe Request* messages on all the available channels. The APs receiving the request will send a *Probe Response* back. When an AP is found whose responses improve the current AP's *Beacons* quality by at least ΔSNR (usually $6 \leq \Delta SNR \leq 8$ dB), the STA initiates a cell switch (reassociation). If a better candidate is not found, the STA returns to the current AP's channel and the scan sweep is repeated periodically.

2.2.3 Security

The design of wired LANs assumes the physical attributes of wires that provide security to the network: authentication is provided by the requirement of a physical connection to the media; confidentiality by the "closed" nature of the wire media. These assumptions are no longer valid for IEEE 802.11 WLANs due to their physically open medium nature. For this reason, the IEEE 802.11 standard [3] defined two services in order to meet the aforementioned assumptions inherited from the wired LANs: authentication and data confidentiality at layer 2. This security is known as Wired Equivalent Privacy (WEP).

The current definition of the standard [24] includes the amendments of 802.11i [18]. IEEE 802.11i improved the security of WEP's authentication and data confidentiality, and also provided data integrity and origin authentication.

Authentication is used to establish the STA's identities before the association process takes place. Two authentication methods are supported: Open System and Shared Key. Open system authentication simply consists of two communications. An *Authentication Request* by the client is followed by an *Authentication Response* from the AP containing a success or failure message.

Shared key authentication relies on the fact that both stations taking part in the authentication process have the same “shared” key, which has been previously set on both STA and AP.

Data confidentiality protects the contents of messages. The current standard [24] provides three cryptographic algorithms: WEP, Temporal Key Integrity Protocol (TKIP) and Counter mode with Cipher-block chaining Message authentication code Protocol (CCMP). By default all messages are sent unprotected.

The original definition of the security services using WEP was found to have severe security weaknesses [74] [120], so a specification of a replacement for WEP became a TGe work item. Later (year 2001) the TGe was split into TGe, devoted to QoS issues, and TGi for security. This group released a new amendment to the standard by the year 2004, known as IEEE4 802.11i. Since then, two levels of security are distinguished: Robust Security Network Association (RSNA) for 802.11i enabled security and pre-RSNA for the old WEP scheme.

2.2.4 Spectrum Management

In order to satisfy requirements in different regulatory domains, specially for the operation of WLANs in the 5 GHz ISM band, a new amendment was released by 2003: the IEEE 802.11h [14]. Two services were specified: Transmit Power Control (TPC) and Dynamic Frequency Selection (DFS).

The TPC service is intended to reduce interference with satellite devices, and DFS is used to avoid co-channel operation with radar systems as well as to ensure uniform utilization of available frequency channels. More details on these services are given in Section 7.1.1.

2.3 Medium Access Control

All IEEE 802.11 STAs in a system share a common wireless medium. The MAC Layer defines the rules that all STAs must follow in order to transmit to this medium. This is handled by using several access mechanisms. The IEEE 802.11 standards define two access methods: Distributed Coordination Function (DCF) and Point Coordination Function (PCF). Moreover, as of late 2005, the IEEE 802.11 task group “E” released a new standard [22] that defines a set of Quality of Service (QoS) enhancements for the IEEE WLANs; it defines procedures for managing network QoS using classes of service. The extensions introduced inherit from the two previous access mechanisms (DCF and PCF). The new MAC protocol is called Hybrid Coordination Function (HCF) (briefly described in Section 2.5), and is only present in QoS-enabled STAs, which are still uncommon nowadays.

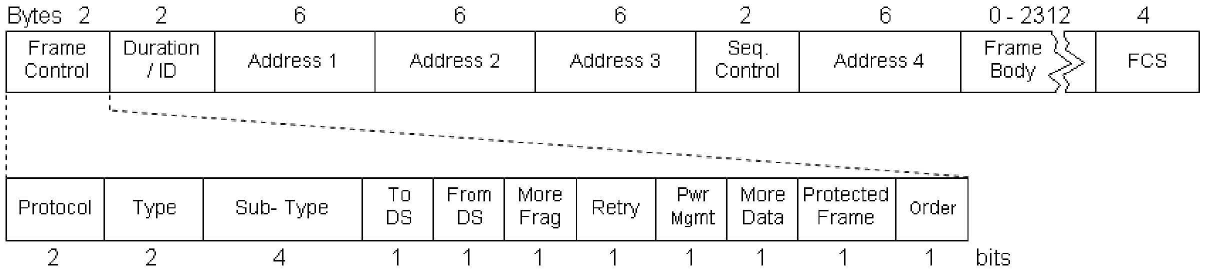

Figure 2.2: Generic 802.11 MAC frame format

Both PCF and HCF are provided through the services of DCF, which is the only mandatory access method. The fundamental access method used by the DCF based MAC is Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). This method is based on contention, whereas the PCF offers a contention free access. The three modes can be used alternately in time.
The optional PCF mode is only allowed in the infrastructure BSS. In this case, the AP polls its associated STAs one after another by sending polling messages. Moreover, if the AP has data ready to be sent to a STA being polled, it can be included in the polling message. If the STA station has data to send to the AP, it is transmitted in the response message. In other words, PCF is a contention free protocol and enables stations to transmit data frames synchronously, with regular time delays between data frame transmissions. The contention free period takes priority over the regular DCF procedure. In this way, delay sensitive packets (e.g. voice or video) can have a higher priority. However, the 802.11 standard is vague in defining portions of the PCF protocol. As a result, APs implementing PCF are rare, even though some chipsets have PCF functionality embedded in the firmware [145]. Moreover, the Wi-Fi Alliance does not include PCF functionality in its interoperability standard.

For all these reasons, in this section we concentrate in the typical DCF mechanism. This section also describes the standard MAC frame format and the carrier sense mechanism.

2.3.1 MAC Frame Format

The IEEE 802.11 MAC defines three types of messages: *data*, *management* and *control*. Management messages are used to support the services described in 2.2. Control frames are used to support the delivery of data and management messages. Data frames carry higher layer information. The general MAC frame format is depicted in Fig. 2.2. Each frame consists of a MAC header, a variable length frame body and an Frame Check Sequence (FCS), which contains an IEEE 32-bit CRC. The MAC header comprises several fields:

- *Frame Control*: consists of several subfields and flags that contain relevant information for STAs in power save mode. It also includes information about the frame type, fragmentation, etc.
- *Duration/ID*: vary with frame type, but it is usually set to the time (in μs) required to complete the current transmission.
- *Address Fields*: used to identify the BSSID, the source address, the destination address, the transmitting STA address and receiving STA address. The contents of the address fields are dependent upon the values of the *To DS* and *From DS* flags in the *Frame Control* field.
- *Sequence Control*: used to number an MAC Service Data Unit (MSDU) and to identify fragments of a given MSDU.
- *QoS Control*: QoS-enabled STAs also include a 2-Byte field to determine the QoS policies desired for the corresponding frame.

Note that not all frames contain all the fields described above. The first three fields and the FCS constitute the minimal format and are present in all frames. The rest of the fields are present only in certain frame types and subtypes. For example, *Address 4* is only present when WDS is used.

2.3.2 Carrier Sense mechanism

Physical and virtual Carrier Sense (CS) functions can be used to determine the state of the wireless medium, which can be either idle or busy. Virtual carrier sense is referred to as the Network Allocation Vector (NAV). The NAV maintains a prediction of future traffic on the medium based on the duration information announced in some frames.

The physical CS functions are provided by the PHY and logically depend on the PHY used. Basically, the physical layer provides a busy/idle medium recognition based on the detection of any energy above a given threshold P_{th}; the physical layer can also report a busy medium by detecting an 802.11 signal (above or below P_{th}). In this case, the physical will consider the medium busy during the time indicated in the PLCP header (see 2.4.1.1 and 2.4.2.2).

The values of P_{th} depend on the modulation used and the transmission power. These values range from -82 dBm (OFDM signals) to -70 dBm (DSSS and transmitted power $\leq$ 50 mW).

2.3.3 DCF based MAC

As mentioned before, the fundamental access method in IEEE 802.11 WLANs is a DCF based on CSMA/CA and a random backoff time following a busy medium condition. In

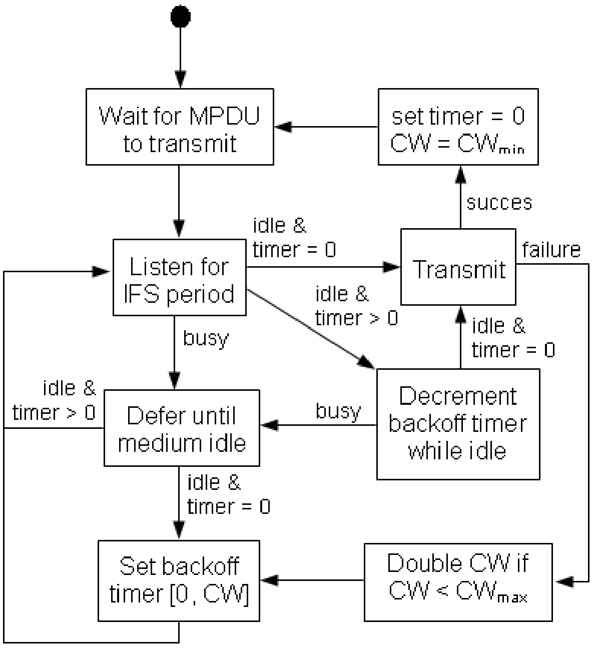

Figure 2.3: DCF MAC procedure

addition, all unicast transmissions use positive Acknowledgment (ACK) frames. If no ACK is received, a retransmission is scheduled by the sender.

The nodes of a WLAN share the medium in a similar manner to the nodes of an Ethernet segment. With CSMA/CA, the nodes sense the air interface before transmitting a frame; if it is busy, they wait until it is released. Physical and virtual carrier sense functions are used to determine the state of the medium (see Section 2.3.2). The CSMA/CA scheme is designed to reduce collision probability when multiple STAs are sharing the same channel. The point where a collision is more likely to occur is just after the medium becomes idle following a transmission. This is because more than one STA may be waiting for the medium to become idle in order to transmit. These potential conflicts are minimized using a random backoff procedure.

2.3.3.1 Basic Access

The DCF's basic access follows the procedure depicted in Fig. 2.3. It works as follows: before initiating a transmission, a station senses the channel to determine whether it is busy (see 2.3.2). If the medium is sensed idle during a period of time called DCF Interframe Space (DIFS), the station is allowed to transmit. If the medium is sensed busy, the transmission is delayed until the channel is idle again.
A slotted binary exponential Backoff (BO) interval is uniformly chosen so that $BO \in [0, CW-1]$, where CW is the contention window. This backoff procedure is invoked before

transmitting a frame when the medium was sensed busy, and also when a transmitting STA infers a failed transmission. A STA uses the CS mechanism to determine whether there is activity during each backoff slot. The backoff timer is decreased as long as the channel is sensed idle, paused when a transmission is in progress, and reactivated when the channel is sensed idle again for more than DIFS. When the backoff timer expires, the station starts transmitting. The value of CW is set to its minimum value, CW_{min}, in the first transmission attempt, and increases integer powers of two at each retransmission, up to a pre-determined value CW_{max}.

After each successfully received unicast data frame, the receiver waits for a Short Inter-frame Space (SIFS) period and transmits an ACK frame (see Fig. 2.4). Logically SIFS $<$ DIFS, thus the ACK has priority over other frames, avoiding unnecessary retransmissions. ACKs are not sent as a response to broadcast/multicast frames.

2.3.3.2 Four-way handshake (RTS/CTS)

The protocol described in the previous subsection is called *basic* or *two-way* handshaking mechanism. In addition, the specification also contains a four-way frame exchange protocol called the RTS/CTS mechanism, which works as follows: a station gains channel access through the contention process described previously, and sends a special frame called Request to Send (RTS), instead of the actual data frame. In response, the receiver sends a Clear to Send (CTS) frame after a SIFS interval. Subsequently, the requesting station is allowed to start the data frame transmission after a SIFS period. See Fig. 2.4 for the complete message exchange sequence for basic access (used by Tx1) and RTS/CTS (used by Tx2).
The information contained in RTS and CTS messages allow all STAs in range to set their NAV to indicate a busy medium until the end of the corresponding ACK transmission.

The main objective of the RTS/CTS handshake is the resolution of the hidden terminal problem [276]. The mechanism is also employed to minimize the lost periods caused by collisions - the RTS frame is much shorter than data fames. Due to the added overhead caused by the extra signaling, this access method is usually deactivated or used only to secure the transmission of long frames (whose size exceeds a given threshold).

2.4 Physical Layer

The IEEE 802.11 standard defines four different transmission techniques with four different PHY implementations: Direct Sequence Spread Spectrum (DSSS), Orthogonal Frequency-Division Multiplexing (OFDM), Frequency-Hopping Spread Spectrum (FHSS) and Infra Red (IR). Note that diffuse infrared and FHSS have received little attention in the market. For this reason, we will focus on the more popular interfaces based on DSSS-CCK (Complementary Code Keying) and OFDM, used under the IEEE 802.11a/b/g specifications.

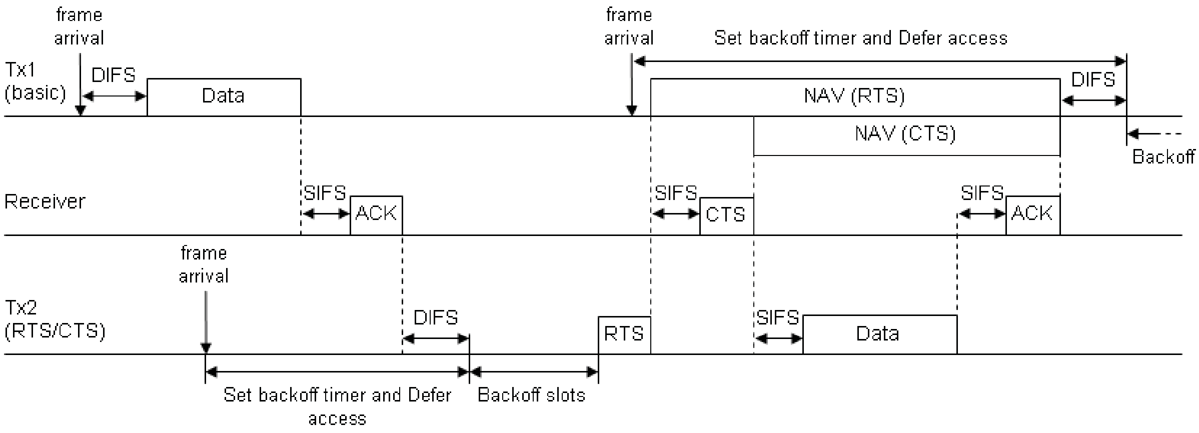

Figure 2.4: Basic and RTS/CTS access with NAV setting

The physical layer is divided into two sublayers: Physical Layer Convergence Procedure (PLCP) and Physical Medium Dependent (PMD). The PLCP sublayer minimizes the dependence of the MAC layer on the PMD sublayer by mapping MAC Protocol Data Unit (MPDU)s into a frame format suitable for transmission by the PMD. Under the direction of the PLCP, the PMD provides actual transmission and reception of PHY entities between two STAs through the wireless medium. To provide this service, the PMD interfaces directly with the air medium and provides modulation and demodulation of the frame transmissions. The PLCP and PMD communicate using service primitives to govern the transmission and reception functions. Note that each PMD sublayer may require the definition of a unique PLCP.

2.4.1 DSSS

This section provides a review of DSSS techniques employed by IEEE 802.11 in the 2.4GHz ISM band. The first specification of the standard defined two basic rates, 1 and 2 Mbps. Later, the 802.11b extension added higher rates (5.5 and 11 Mbps) as a natural extension of legacy DSSS by using a Complementary Code Keying (CCK) scheme. This new scheme provided a smooth transition to higher rates with DSSS, and it also allowed interoperability with the old 1 and 2 Mbps devices.

2.4.1.1 PLCP Frame Format

The 802.11 frame structure is shown in Fig. 2.5. The PLCP preamble is formed from a *SYNC* field and an *Start Frame Delimiter (SFD).* The *SYNC* field consists of scrambled 1 bits so that the receiver can perform the necessary synchronization operations (e.g. signal detection, timing acquisition, multipath estimation, etc.). The *SFD* indicates the start of PHY-dependent parameters. The 48 bit PLCP header contains the overhead needed

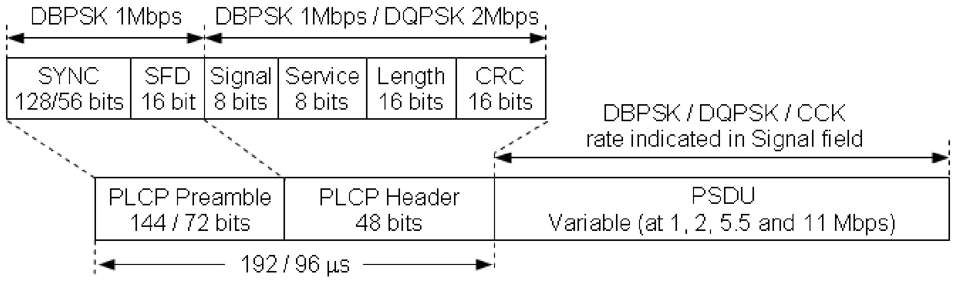
DBPSK 1Mbps
DBPSK 1Mbps / DQPSK 2Mbps
SYNC 128/56 bits
SFD 16 bit
Signal 8 bits
Service 8 bits
Length 16 bits
CRC 16 bits
DBPSK / DQPSK / CCK rate indicated in Signal field
PLCP Preamble 144 / 72 bits
PLCP Header 48 bits
PSDU Variable (at 1, 2, 5.5 and 11 Mbps)
192 / 96 µs

Figure 2.5: PPDU format for DSSS/CCK 802.11 PHY

by the PHY. The IEEE 802.11b standard [5] defines two different frame structures: *long* and *short* preambles. The long preamble consists of a 128 bit *SYNC* field, 16 bits of *SFD* and PLCP header sent at the lowest rate, 1 Mbps. Then, the duration of a *long preamble*, including PLCP header, is 192 μs. On the other hand, the short preamble option was intended to reduce the overhead of transmissions. The short preamble reduces the *SYNC* field by half and increases to 2 Mbps the rate used to send the PLCP header. In this way, the duration of a *short preamble* is 96 μs.

2.4.1.2 Legacy 802.11 DSSS

DSSS spreads a transmitted signal across a wider spectrum range in order to provide a spreading gain against narrowband interference. DSSS techniques spread transmissions by XORing a stream of data bits with a pseudo-random sequence. In 802.11 [3], this sequence of bits is always the 11 bit Barker sequence: $+1-1+1+1-1+1+1+1-1-1-1$. That is to say, each data symbol is transmitted as 11 bits represented by the original Barker sequence or its inverse. Note that for all available data rates (from 1 to 11 Mbps), the chip rate is 11 Mcps, resulting in a channel bandwidth of 22 MHz.

For 1 Mbps, the Barker sequence is used to spread a 1 Mbps Differential BPSK (DBPSK) signal. For 2 Mbps, the same sequence is Differential QPSK (DQPSK) modulated. With differential PSK modulations, the information is conveyed by the phase difference between adjacent transmitted signal elements. In this way, the system complexity is reduced, since a coherent phase reference is not needed. On the other hand, the Bit Error Ratio (BER) is higher for a given Signal-to-Noise Ratio (SNR). The raw baseband data is differentially encoded as follows: if x_i is the bit intended for transmission and y_i is the bit actually transmitted (differentially encoded); $y_i = y_{i-1} \oplus x_i$.

2.4.1.3 High Rate DSSS

In 1999, the IEEE 802.11b [5] was released. This extension defined two additional PHY rates of 5.5 and 11 Mbps and is known as High Rate DSSS (HR-DSSS). However, the data is no longer spread using the Barker sequence. Instead, CCK is used, and optionally, Packet Binary Convolutional Coding (PBCC). Since support for PBCC is optional while CCK is mandatory, and no other vendor, apart from Texas Instruments seems interested in PBCC, this section will focus exclusively on CCK.

The same 802.11 PLCP header based on 1 Mbps is maintained. Moreover, to insure that the modulation has the same bandwidth and waveform as the existing 802.11 DSSS modulation, the chipping rate is kept at 11 Mcps. Therefore, this approach enables back-compatibility with the older 802.11. The symbol rate is increased from 1 to 1.375 MSps; there are 8 bits/symbol at 11 Mbps and 4 bits/symbol at 5.5 Mbps.

CCK is a variation on M-ary Orthogonal Keying [50] developed by Lucent Technologies and Harris Semiconductor. It utilizes a set of Complementary Codes that are derived from incoming data.
For 11 Mbps, each group of 8 bits of incoming data is split into 4 dibits. The first dibit is converted to a phase angle φ_1, which is differentially encoded across successive code words (similar to the 1 and 2 Mbps schemes). The rest of the phases (φ_2 to φ_4) define a set of 64 possible codes of 8 chips and φ_1 can be thought of as an extra phase rotation to the entire code word (see eq. 2.1). Please refer to [152] for more information on these codes. The 8 complex chips are mapped to the same Quadrature Phase-Shift Keying (QPSK) constellation and transmitted.
This process can be explained from a different perspective. We can think of each chip as a function of φ_2, φ_3 and φ_4, then, instead of encoding the chips with QPSK, the entire symbol is encoded using DQPSK based on φ_1 (phase difference with the previous symbol) [159]. That is to say, 6 information bits select 1 of 64 possible codes and 2 select 1 of 4 DQPSK phases.

$$\begin{aligned} c = \{&e^{j(\varphi_1+\varphi_2+\varphi_3+\varphi_4)}, e^{j(\varphi_1+\varphi_3+\varphi_4)}, e^{j(\varphi_1+\varphi_2+\varphi_4)}, -e^{j(\varphi_1+\varphi_4)}, \\ &e^{j(\varphi_1+\varphi_2+\varphi_3)}, e^{j(\varphi_1+\varphi_3)}, -e^{j(\varphi_1+\varphi_2)}, e^{j(\varphi_1)}\} \end{aligned} \tag{2.1}$$

The 5.5 Mbps version is carried out similarly. In this case, only 4 bits are used to find the 8 chip CCK code. The first two bits are used to find φ_1 (phase rotation for the whole code word) in exactly the same manner as with 11 Mbps. The other two bits are used to select 1 of 4 possible codes. The chip rate and symbol rate are still 11 Mcps and 1.375 MSps, respectively, and in consequence channel bandwidth and waveform remain the same.

2.4.2 OFDM

Orthogonal Frequency-Division Multiplexing (OFDM) is the basis for IEEE 802.11a [4] and 802.11g [13] standards that operate in the 5 and 2.4 GHz ISM bands, respectively. OFDM along with HR-DSSS, are the most popular technologies for WLAN radios. This section provides a brief review of these techniques.

Mode	PHY Rate (Mbps)	Modulation	Coding rate	Data bits per Symbol
m1	6	BPSK	1/2	24
m2	9	BPSK	3/4	36
m3	12	QPSK	1/2	48
m4	18	QPSK	3/4	72
m5	24	16-QAM	1/2	96
m6	36	16-QAM	3/4	144
m7	48	64-QAM	2/3	192
m8	54	64-QAM	3/4	216

Table 2.2: Modes of the OFDM 802.11 PHY

2.4.2.1 Modes

With OFDM, the high rate data stream is split into 52 lower rate streams that are transmitted simultaneously in different sub-carriers. Thus, the effects of time dispersion caused by multi-path propagation are reduced. Four of the sub-carriers are pilot sub-carriers that are used as a reference to disregard frequency or phase shifts. Pilots are scrambled by a length 127 pseudo-noise sequence. The remaining 48 sub-carriers provide separate wireless “pathways” for sending the information in a parallel fashion.

The standards define eight different modes providing 6 to 54 Mbps. For all these modes, the symbol duration is 4 μs. Moreover, in order to reduce intersymbol interference, a guard time of 800 ns is introduced in every OFDM symbol. The guard time is a repeat of the end of the symbol at the beginning. This technique, known as cyclic prefixing, is used to retain sinusoid properties that ease channel estimation in multi-path channels. The symbol duration and guard time determine sub-carrier spacing to be 0.3125 MHz ($(4\mu s - 800ns)^{-1}$) and also the total bandwidth used, 16.6 MHz for all modes.
The slowest mode uses Binary Phase-Shift Keying (BPSK) and the fastest uses 64-Quadrature Amplitude Modulation (QAM). With 48 sub-carriers, 4 μs symbols and 1 to 6 bits per sub-carrier, the raw bit rate ranges from 12 to 72 Mbps. To correct for sub-carriers in fades, forward error correction is applied by means of convolutional codes. Coding rates are variable (see Table 2.2) and hence the eight different admitted modes are: 6, 9, 12, 18, 24, 36, 48 and 54 Mbps (only 6, 12 and 24 Mbps are mandatory).

2.4.2.2 PLCP Frame Format

Similar to the DSSS PHY, OFDM also requires a preamble preceding every frame. This preamble and the PLCP header are essential to successfully decode actual data bits.

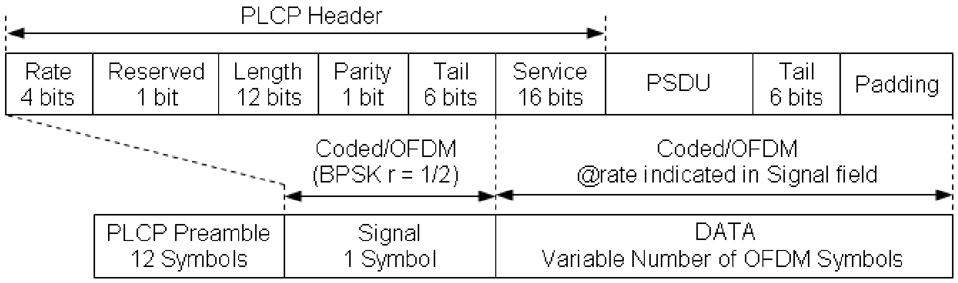

Figure 2.6: PPDU format for OFDM 802.11 PHY

Figure 2.6 shows the structure of a PLCP Protocol Data Unit (PPDU). The PPDU includes PLCP preamble, PLCP header, MPDU (from MAC Layer), tail bits, and pad bits, if necessary.

The first part of the preamble consists of 10 repeated symbols, the duration of which is 800 ns. These short symbols allow a coarse frequency offset estimation and provide a convenient way of performing frame detection and automatic gain control [246]. Subsequently, there are two long training symbols that contain 52 QPSK modulated subcarriers (like a normal symbol), adding 8 μs to the preamble. The long symbols are used to obtain reference amplitudes and phases for a coherent demodulation. The PLCP header, except for the *Service* field, contributes a single OFDM symbol (4 μs), which is transmitted with BPSK modulation and 1/2 coding rate. The Tail bits are set to '0' to return the convolutional code to the "zero" state. The pad bits are used to make the resulting bit string into a multiple of OFDM symbols. The portion of the PPDU represented by *DATA* field (Fig. 2.6) is transmitted at the data rate specified in the *Rate* field.

2.4.2.3 OFDM in the 2.4 GHz band

In 2003, the IEEE 802.11g [13] amendment was released to provide higher data rates in the 2.4 GHz band used by the DSSS specifications. The new PHY is known as Extended Rate PHY (ERP). The modulation scheme used in 802.11g is copied from 802.11a [4], but reverts to CCK (like 802.11b [5]) for 5.5 and 11 Mbps and DSSS for 1 and 2 Mbps. Thus 802.11g hardware is fully backwards compatible with older 802.11b hardware. However, there are two additional optional modulations defined: ERP-PBCC (22 and 33 Mbps) and DSSS-OFDM. The first is an extension of the PBCC optional modulation mentioned in 2.4.1.3, and the latter is a hybrid modulation combining a DSSS preamble and header with an OFDM payload transmission.

Consequently, an ERP STA will support at least three different preamble and header formats: the long and short preambles defined in 2.4.1.1 and the OFDM preamble described in 2.4.2.2.

2.5 Alphabet Soup

The IEEE released the initial standard for wireless LANs, IEEE 802.11 [3] in June 1997. This standard specified PHY and MAC layers for operation in the 2.4 GHz ISM band with data rates of 1 and 2 Mbps. As mentioned earlier, the initial version of 802.11 defined FHSS or DSSS. Since the ratification of the initial standard, the IEEE 802.11 WG has made several revisions through various task groups.

The terms "standard" and "amendment" are sometimes used interchangeably when referring to the different flavors of 802.11, but as far as the IEEE Standards Association is concerned, there is only one current standard; it is denoted by IEEE 802.11 followed by the date it was published. The standard is updated by means of amendments. These amendments are created by different task groups (TG). Task groups within the 802.11 WG [1] enhance portions of the 802.11 standard. A particular lower-case letter corresponding to each standard/revision is used to identify the different task groups. For example IEEE 802.11a, 802.11b, 802.11g, etc. The most recent version of the standard (at the time of writing) is denoted as IEEE 802.11-2007 [24] and it includes the amendments defined by the task groups IEEE 802.11a, b, d, e, g, h, i and j.
Updating 802.11 is the responsibility of TGm. In order to create a new version, TGm combines the previous version of the standard and all published amendments. TGm also provides clarification and interpretation to industry on published documents.

In this section we give an overview on some of the amendments. Some of them are already in use, some are nearing completion and others are just starting up.

802.11a: OFDM in the 5 GHz ISM band This amendment is now part of the standard and has already been explained in previous sections (see 2.4.2). It defines an OFDM-based PHY offering 54 Mbps in the 5 GHz ISM band. This band offers much less potential for radio frequency interference than other PHYs (e.g., 802.11b and 802.11g) that utilize the 2.4 GHz band. On the other hand, a higher working frequency implies reduced coverage, since propagation loss is directly proportional to the frequency. It was ratified in 1999, although 802.11a-based products became available in late 2001.

802.11b: High Rate DSSS in the 2.4 GHz ISM band TGb was responsible for enhancing the initial DSSS PHY to allow 5.5 and 11 Mbps by means of CCK coding, as detailed in Section 2.4.1.3. It was released in 1999 [5] and modified in 2001 [8]. Most wireless LAN installations today comply with 802.11b, which is also the basis for Wi-Fi certification[2].

802.11c: Bridge operation This is actually a supplement to the IEEE 802.1D MAC bridges standard. Ratified in October 1998, it adds requirements associated with bridging 802.11 wireless client devices.

[1] http://www.ieee802.org/11/

[2] http://www.wi-fi.org

802.11d: Additional regulatory domains To date, only a few regulatory domains have rules in place for the operation of 802.11 wireless LANs. This supplement, released in 2001, adds the requirements and definitions necessary to allow 802.11 WLAN equipment to operate in markets not served by the current standard. More precisely the 802.11d allows configurations at the MAC level intended to comply with the rules of the country or district in which the network is to be used. Rules subject to variation include allowed frequencies, allowed power levels, and allowed signal bandwidth. Hence, the 802.11d eases global roaming: the same device can be used abroad, regardless of the local regulations.

802.11e: QoS enhancements As of late 2005, the IEEE 802.11 TGe released a new amendment that defines a set of QoS enhancements for the IEEE WLANs [22]; it defines procedures for managing network QoS using classes of service. The extensions introduced consider the two access mechanisms: DCF and PCF. The new MAC protocol is called Hybrid Coordination Function (HCF). It combines a contention channel access mechanism, the Enhanced DCF Channel Access (EDCA), and a polling-based channel access, the HCF Controlled Channel Access (HCCA).

EDCA is designed to manage 8 different traffic priorities. Packets belonging to the different traffic priorities are mapped into 4 access categories (ACs) buffered in separate queues at a station - background, best effort, video and voice; each of them represents a different priority level defined by different AC-specific CW sizes, Arbitration IFS (AIFS) values and Transmission Opportunity (TXOP) limits. Each AC within a station contends for the channel independently of the others according to a CSMA/CA access, but in this case, an internal (virtual) collision between two ACs within a station is handled by granting the access to the AC with the highest priority. The higher priority ACs are assigned smaller AIFS (i.e. higher priority ACs enjoy a relatively faster progress through backoff slots) and may select backoff values from a comparably smaller CW range (i.e. higher priority ACs have smaller backoff delays). Upon gaining the access to the medium, each AC may carry out multiple frame exchange sequences as long as the total access duration does not go over a TXOP limit.

The HCCA (which is not mandatory) works like the PCF but adding traffic classes. HCCA is the most complex coordination function. With HCCA, QoS can be parameterized with great precision. However, IEEE 802.11e enabled devices usually implement EDCA only, which is not intended to offer any form of QoS guarantee to the users in terms of bandwidth allocation, bounded delay, etc. Nonetheless, this service prioritization is translated into an increment of the probability to win the contention for the access to the common channel.

In addition to HCCA and EDCA, 802.11e specifies additional optional protocols for enhanced 802.11 MAC layer QoS: Automatic Power Save Delivery (APSD) is a more efficient power management method than legacy 802.11 Power Save Polling; Block Acknowledgments allow an entire TXOP to be acknowledged in a single frame; Direct Link Setup (DLS) allows direct station-to-station frame transfer within a BSS.

802.11F: Inter Access Point Protocol (IAPP) A station's handoff calls for AP to AP communication, but this process was not defined in the original 802.11 specification. Each vendor had its own particular implementation, and the publication of the IAPP [11] was intended to facilitate interoperability. The amendment was released as a Trial Use Recommended Practice in 2003. The protocol provided the necessary capabilities to achieve multi-vendor AP interoperability across a DS supporting IEEE P802.11 WLAN links. It was designed for the enforcement of unique association throughout an ESS, and for secure exchange of the station's security context between the current AP and the new AP produced as a result of the station's roaming.
IAPP calls for the new servicing AP to send out two packets onto the DS. One of them is actually set with the client's source address and is used to update switches' MAC address tables with the client's new location. The other is sent to an IAPP multicast address that all APs subscribe to, and contains the MAC address of the new associated station. All APs will receive this packet, and the one that had been associated with that station will remove the stale association from its internal table. Furthermore, IAPP provides for the sharing of generic information between APs by means of some not well defined "contexts". This lack of specification made IAPP of little practical interest as far as vendor interoperability is concerned. For these reasons this standard practice has been withdrawn as of February 2006.

802.11g: OFDM in the 2.4 GHz ISM band This amendment is now part of the standard and has already been explained in previous sections (see 2.4.2). It defines an OFDM-based PHY offering 54 Mbps in the 2.4 GHz ISM band [13]. It was ratified in 2003, although pre-standard 802.11g-based products became available months before. Its definition is identical to the 802.11a core, except for some additional legacy overhead for backward compatibility: 802.11g implements all mandatory elements of the IEEE 802.11b PHY standard. In spite of the TG efforts to enable interoperability and coexistence of 11b and g devices, the presence of a legacy 802.11b participant will significantly reduce the performance of the overall 802.11g network.

802.11h: Spectrum managed Published in October 2003, the IEEE 802.11h [14] was intended to resolve problems arising from the coexistence of 802.11 networks with radar and satellite systems in the 5 GHz band (i.e. 802.11a). It was originally designed to address European regulations, but is now applicable in many other countries. As detailed in Section 7.1.1, the standard provides Dynamic Frequency Selection (DFS) and Transmit Power Control (TPC) capabilities to IEEE 802.11a WLANs.
DFS is a required functionality for WLAN operating in the 5 GHz band to avoid co-channel interference with radar systems and to ensure uniform utilization of available channels.

TPC is required to ensure that all devices in a 802.11a WLAN comply with the transmission power limitations applicable to the channel in use in a given region so that the possible interference with satellite/radar systems is minimized.

802.11i: Enhanced security The IEEE 802.11i amendment, released in July 2004, specifies strong security mechanisms for IEEE WLANs. This amendment was integrated into the current version of the standard [24]. The new security specifications provided by TGi supersede the previous WEP, which was shown to have severe security weaknesses [74] [120]. The Wi-Fi Alliance responded to that weakness by introducing Wi-Fi Protected Access (WPA), which implemented a subset of the 802.11i specification. The full implementation of the IEEE 802.11i is also known as WPA2 (according to the Wi-Fi Alliance nomenclature) and provides Robust Security Network (RSN).

WPA is based on TKIP, a set of algorithms that improved WEP security without the need of major upgrades, since it also uses RC4 as its cipher. TKIP improves confidentiality by implementing key mixing and re-keying functions. Moreover, it implements a message sequence counter and 64-bit check known as MICHAEL to add integrity and resilience to reply attacks. WPA also enables IEEE 802.1X based authentication.

The IEEE 802.11i (and WPA2) introduced the concept of RSN, which included the improvements of WPA but adding CCMP (used instead of TKIP), based on the block cipher Advanced Encryption Standard (AES) (instead of the RC4 stream cipher). While this combination provides confidentiality, integrity and origin authentication, it is not compatible with the old WEP-based hardware since it requires a dedicated chip to handle the encryption and decryption.

802.11j: Extensions for operation in Japan Finalized in 2004, this amendment to the IEEE 802.11 standard is designed specially for the Japanese market. The idea is to obtain Japanese regulatory approval by enhancing the 802.11 MAC and 802.11a PHY to additionally operate in newly available Japanese 4.9 and 5 GHz bands for indoor, outdoor and mobile applications.

802.11k: Radio measurements The work of the TGk group (Radio Resource Measurement) resulted in the IEEE 802.11k amendment [36], released in December 2008. It is intended to improve the provision of data from the physical and medium access layers by defining a series of measurement requests and reports that can be used in upper layers. With the improvements introduced by the new standard, measurements are not only carried locally, but it will also be possible to request measurements from neighboring stations. In order to grant the possibility of making measurement requests and to inform other stations about measurements, 802.11k inherits several mechanisms that were already standardized by the IEEE 802.11h.

The resulting information is made available to upper layers for any purpose, being especially useful for radio resource management strategies. That is to say, 802.11k only provides useful information on several radio parameters; the design of radio resource management mechanisms is out of its scope. More details are given in Section 7.1.2.

802.11n: Further throughput improvements TGn is currently investigating new improvements to the 802.11 PHY and MAC layers to provide high throughput (near 600 Mbps). The work on the 802.11n amendment was initiated in 2004. Its definition is expected to be finalized by December 2009, but major manufacturers have been releasing pre-standard products based on early specifications. Basically those improvements are obtained by adding multiple-input multiple-output (MIMO) techniques and channel bonding to the PHY layer, and frame aggregation to the MAC layer.

Aggregation is a process of packing multiple MSDUs or MPDUs together to reduce the overheads and average them over multiple frames. The aggregation requires the use of Block Acknowledgment (BlockAck), which was previously introduced in 802.11e.

Channel Bonding, refers to the ability of using two separate non-overlapping channels simultaneously to transmit data. This mode of operation doubles the occupied bandwidth, and hence the PHY rate of a “traditional” 20 MHz channel is effectively doubled.

MIMO uses multiple antennas to coherently resolve more information than possible using a single antenna. These techniques provide two main benefits to 802.11n: antenna diversity and spatial multiplexing. First, MIMO uses the multipath signal diversity to increase a receiver’s ability to recover the message information from the signal. And second, the spatial multiplexing can significantly increase data throughput by sending independent data streams that are transferred simultaneously within one spectral channel of bandwidth. The number of simultaneous data streams is limited by the minimum number of antennas and RF chains in use on both sides of the link.

802.11p: Vehicular networks TGp will define enhancements to the 802.11 standard to support Intelligent Transportation Systems (ITS) applications. The new amendment, also known as Wireless Access in the Vehicular Environment (WAVE), will provide support for data exchange between high-speed vehicles and communications between vehicles and roadside infrastructure. Communications may be possible at speeds of up to 200 km/h and within a range of 1 Km. Very short latencies will also be needed as some applications must guarantee data delivery within few milliseconds. IEEE 802.11p will also be applicable to marine and rail communications. The most evident applications of such a vehicle communication network include vehicle safety services, traffic jam alarms, toll collections, vehicle collision avoidance, and adaptive traffic light control.

802.11r: Fast roaming Published in 2008, the IEEE 802.11r enhancements are intended to minimize or eliminate the connectivity disruption periods experienced by users during a BSS transition, i.e. during a handover between two APs of the same ESS. The idea behind these improvements was to support real time constraints imposed by applications, such as VoIP. New QoS and security mechanisms introduced by 802.11e and 11i amendments added an extra burden to the handoff process, which was initially much simpler (just a 4-frame exchange). The addition of those new features increased the number of required messages dramatically. During the time these additional messages are being exchanged, the mobile device's traffic cannot proceed, and the user will notice connectivity loss of the order of few seconds. The Fast BSS Transition defined by 802.11r eliminates much of the handshaking overhead by allowing both the security key negotiation and requests for wireless resources to occur in parallel before the AP transition is actually performed.

802.11s: Mesh networking The purpose of the TGs is to provide a protocol for auto-configuring paths between APs over self-configuring multi-hop topologies in a WDS to support both multicast and unicast traffic in an ESS mesh. A device in an 802.11s mesh network is labeled as a Mesh Point (MP). These form mesh links with one another, over which mesh paths can be established by means of the Hybrid Wireless Mesh Protocol (HWMP), a routing protocol inspired by a combination of Ad hoc On-Demand Distance Vector (AODV) and tree-based routing. An MP can also perform as a regular AP in order to provide access to the mesh network to legacy 802.11 clients.

While there are still ongoing efforts to improve the IEEE 802.11s Draft, a reference implementation of the 802.11s is available as part of the mac80211 layer in the Linux kernel, starting with version 2.6.26.

802.11v: Network management TGv is working (at the time of writing) on an amendment to the 802.11 standard [37] to allow configuration of client devices while connected to IEEE 802.11 networks. It will enable management of attached stations in a centralized (cellular-like) or in a distributed fashion through a layer 2 mechanism. This includes, for example, the network's ability to monitor, configure, and update client STAs. While it extends prior work in radio measurement to effect a complete and coherent upper layer interface for managing 802.11 devices (the TGk defined messages to retrieve information from the station, the ability to configure the station was not within their scope), early reports revive fears of duplication between standards bodies, describing 802.11v as similar in nature to the IETF Control and Provisioning of Wireless Access Points (CAPWAP). However, CAPWAP is limited to the AP management, while 802.11v involves client STAs. IEEE 802.11v amendment is scheduled for ratification in 2010.

Besides improved management, works within TGv are broken down into four other categories: power saving mechanisms with Wi-Fi handheld VoIP devices in mind; positioning, to enable new location-aware services; timing to support applications that require very precise calibration; and coexistence, to reduce co-located interference. More detailes are given in Section 7.1.3.

802.11y: 3650–3700 MHz Operation Released in November 2008, 11y [35] enables 802.11 equipment to operate in the 3650 to 3700 MHz band (except when interferring with a satellite earth station) in USA, although other bands in different regulatory domains are being studied. The Federal Communications Commission (FCC) rules at 3650 MHz allow for registered stations to operate at much higher power than in traditional ISM bands (up to 20 W Equivalent Isotropically Radiated Power (EIRP)). Other three concepts are added: Contention Based Protocol (CBP), Extended Channel Switch Announcement (ECSA), and Dependent Station Enablement (DSE). CBP includes enhancements to the carrier sensing and energy detection mechanisms of 802.11. ECSA provides a mechanism for an access point to notify the stations connected to it of its intention to change channels or to change channel bandwidth. Finally, DSE is used to manage licensing issues.

Chapter 3

Characterizing the environment

Before starting a discussion about the optimization of radio resources, it is imperative that we understand the characteristics of the radio channel. The best channel assignment, or transmitted power configuration cannot be obtained without considering the propagation loss, the receiver design, the modulation performance, or the effects of interference. This chapter explains the phenomena of propagation and the effects of noise and interference. For all these phenomena, the study includes a theoretical analysis that is compared with empirical experiments.

Propagation has been widely studied in the literature, and this chapter contributes with a review of the models that better describe propagation in WLAN scenarios. Then, the performance of the available modulations is studied in terms of BER vs. SNR. The formulation for most of the modulations is well known. However, a deeper literature research was needed for the CCK case. Subsequent sections are devoted to studying the effect of interference. A considerably large amount of research papers studying the coexistence between different systems (e.g. WLAN and WPAN) can easily be found. But ironically, however, interference between IEEE 802.11 transmitters has not received the same attention, although its understanding is fundamental for designing a good frequency management strategy. For this reason, this chapter also presents the contributions to this study developed in this thesis [129]. Bearing in mind all these issues, this chapter finally provides a comparison of different simulation environments, paying special attention to their propagation, error and interference models. This study is needed to establish the limitations of our simulations, which are a consequence of the approximations and assumptions made by the simulation environment.

3.1 Propagation

In a radio receiver, the signal power available at the terminals of the antenna depends on the power density of the electromagnetic wave generated by the transmitter. The received power P_r in watts, is written as:

$$P_r = WA \tag{3.1}$$

where A (in m^2) is the antenna effective area, and W the power density of the electromagnetic wave (in W/m^2). W depends on the power supplied at the transmitter antenna (P_t), and the distance d (in m) between transmitter and receiver antennas. A and W are given by:

$$A = \frac{\lambda^2 G_r}{4\pi} \tag{3.2}$$

$$W = \frac{P_t G_t}{4\pi d^2} \tag{3.3}$$

G_t and G_r are the antenna gains for the transmitter and the receiver, with respect to an ideal isotropic antenna radiation pattern. The wavelength of the signal is represented by λ. From the above formulas (3.1, 3.2 and 3.3), we can rewrite P_r as:

$$P_r = P_t G_t G_r \left(\frac{\lambda}{4\pi d}\right)^2 \tag{3.4}$$

which is known as the Friis transmission equation. This formula expresses the relationship between propagation loss, the distance between antennas and the frequency used. However, it is not accurate in many real-life environments, where multipath propagation may occur due to the presence of obstacles, and the possibility of having no direct line of sight between transmitter and receiver. Many empirical models have been proposed in the literature that try to provide more accurate predictions.
Propagation loss, or path loss (L_P), defined as the ratio of radiated to received power, is usually expressed in dB. The following formula represents L_P for f in MHz and d in m, and given that $\lambda = c/f$ (c is the speed of light) and $20log_{10}(c/4\pi) = 27.6dB$:

$$L_P[dB] = 20log_{10}f + 20log_{10}d - 27.6 \tag{3.5}$$

3.1.1 Indoor propagation

From eq. 3.4 it is clearly shown that the propagation loss is proportional to $(1/fd)^\gamma$, where f is the operating frequency and $\gamma = 2$. Even though this is known to be inaccurate for most of the scenarios, it is common practice to use it as a basis for estimates, in which case a value for $\gamma > 2$ is chosen according to the environment [256].
There is a wide range of indoor propagation models due to the great diversity of building profiles. Different buildings may have different building materials, large or small rooms, high or low ceilings, furnishings, etc. and hence different propagation models should be applied accordingly.

A breakpoint model, as described in [246], is the most simple adaptation for indoors. It uses $\gamma = 2$ for the first meters, and for distances larger than a defined breakpoint (e.g.

5 m) γ is increased to 3.5.
In [278] a generalization of other existing models is given, which captures all possible effects on the signal propagation:

$$L_P[dB] = k_1 + G_a + k_2 log_{10} f + k_3 log_{10} d + n_w(k_4\Phi_1 + k_5\Phi_2) + k_6 n_f \tag{3.6}$$

In 3.6, G_a, f y d represent the antenna gains, frequency and distance, respectively. Φ_1 y Φ_2 depend on the angle between walls and incident waves, n_w and n_f are the number of walls and the number of floors that the signal traverses. The coefficients k_i are determined by means of practical measurements. This allows a good adaptation to most kinds of indoor scenarios.
A lighter expression is suggested in [256]:

$$L_P(d) = L_P(d_0) + 10\gamma log_{10}\left(\frac{d}{d_0}\right) + X_\sigma \tag{3.7}$$

where d_0 represents the path loss at a reference distance close to the transmitter (e.g. $d_0 = 1$ m). Different receivers placed at the same distance from one transmitter will receive a different signal due to the different paths traveled by the waves. Furthermore, the signal varies in time as a consequence of a dynamic environment. For these two reasons, eq. 3.7 adds the random variable X_σ, following a log-normal distribution with standard deviation σ. After empirical measurements, it was concluded that σ ranges between 2 and 9 dB, depending on the environment.
This formula can also adapt the distant exponent γ. For an open plan building, the suggested exponent is $\gamma = 2.2$, for a semi-open office profile $\gamma = 3.3$; finally, $\gamma = 4.5$ for closed office (i.e. great density of walls and furniture).

3.1.2 Outdoor propagation

WLAN outdoor applications are characterized by typical antenna heights of a few meters, distance ranges of tens of meters to one or two kilometers and urban, suburban or campus-type scenarios. IEEE 802.11 standards are not intended for long-distance radio links and limit the cell radius to no more than 3 km [232] [185]. However, some technical and legal limitations may be overcame to obtain 40 km links [176], but these issues are beyond the scope of this thesis.

With low antennas and relatively long distances between transmitters and receivers, the ground invades the first Fresnel zone causing attenuation for diffraction, even in a flat terrain scenario. Similar to the indoor breakpoint model, a double breakpoint model is also described in [246] in order to capture the outdoor propagation phenomena. This model has a first breakpoint at 1 m and a second breakpoint at above 100 m.

Various models exist that are derived from well-known empirical approaches made in the field of cellular communications. The Okumura-Hata [154] model is limited to

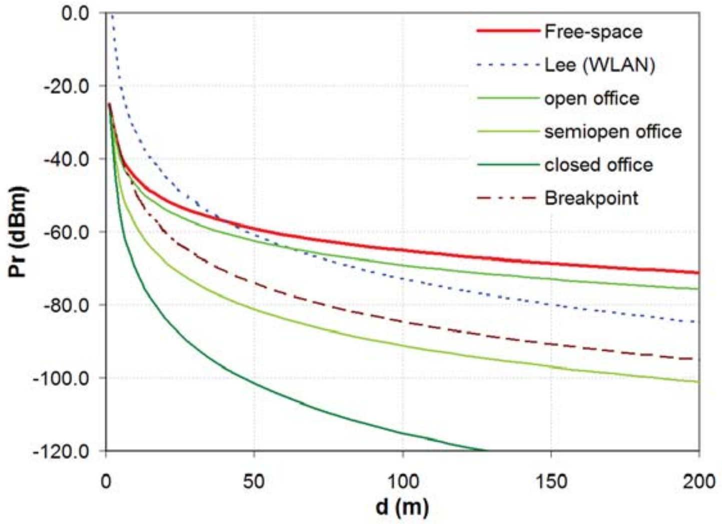

Figure 3.1: Received signal power with different indoor/outdoor propagation models

frequencies under 1500 MHz, antennas placed at an altitude of more than 30 m and, similar to the Longley-Rice [206] model, it is verified for distances from 1 km; the COST231-Hata [230] is an alternative for frequencies around 2 GHz. Lee's loss model [196] is not frequency dependent, but it takes into account the extra gain due to antenna heights and increases by a factor d^4. In [149] Lee's model is adapted for the specific characteristics of an outdoor WLAN. In this adaptation, a frequency factor is introduced in the likeness of the free-space model:

$$L_P[dB] = 40log_{10}d + 20log_{10}f - 20log_{10}h_th_r \tag{3.8}$$

where d is, as usual, the distance between transmitter and receiver, f is the frequency in GHz, the transmitter's antenna height is h_t and the receiver antenna is at h_r m. There are other models in the literature that try to reflect particular scenarios (e.g. see [183] for a campus-type WLAN). Figure 3.1 shows the evolution of received power as the distance d is increased, according to different models (P_t = 15 dBm, d_0 = 1 m, $L_P(d_0)$ = 35 dB and G_x = 0 dB).

3.1.3 Coverage estimations

The term *cell radius* is used in the previous subsection to refer to the longest possible distance between a transmitter and a receiver. In an infrastructure-based WLAN, this distance is measured between the AP and the farthest client STA. The cell radius depends on the transmitted power, the noise and the presence of interference.

The filtered white noise at the receiver is computed as $N = kTB$; where k is the Boltzmann constant ($1.38{\cdot}10^{-23} J/K$), T is the temperature (in Kelvin) and B the receiver fil-

PHY	**DSSS/CCK**				**OFDM**							
PHY rate [Mbps]	1	2	5.5	11	6	9	12	18	24	36	48	54
Rec. sens [dBm]	-93	-90	-87	-84	-82	-81	-79	-77	-74	-70	-66	-65
Cell radius [m]												
Free Space	1380	980	695	495	390	348	276	220	155	98	62	55
Lee (WLAN)	240	202	170	145	128	121	107	96	81	64	51	48
Open office	715	521	382	280	227	204	166	134	98	65	43	38
Semiopen office	85	65	53	45	38	36	31	26	21	16	12	11
Closed office	25	22	19	17	15	14	12	11	10	8	7	6

Table 3.1: Cell radius for different modulations and propagation models

ter's bandwidth. For B=22 MHz, T=295 K, the filtered noise power is $N = -100.48$ dBm. Furthermore, additional system noise is introduced at the early stages of the receiver. This contribution of the RF components to the overall noise is represented by the noise figure F, so that $F = SNR_{in}/SNR_{out} \geq 1$. This value is typically between 4 and 6 dB in most IEEE 802.11 devices.

Besides white thermal noise and system noise, there is a third contribution known as man-made noise [69], different in nature from in-band interference signals. Man-made noise can be produced by different sources (e.g. engines, electrical equipment) and depends on the frequency band under study. According to [246], typical man-made noise levels for the 2.4 GHz band at indoor environments are around 6 dB above thermal noise.

The total receiver noise determines the effective SNR required to allow reliable operation with sufficiently low error rate. Because of the minimum required SNR, there is a requirement for the minimum required P_r (received signal power). The IEEE 802.11 specifies this minimum received P_r, known as receiver sensitivity, as the input level measured at the antenna connector that produces a BER $\leq 10^{-5}$. Therefore, receiver sensitivity indicates how faint an RF signal can be successfully received by the receiver. The IEEE 802.11 standards [4, 5, 13] define the minimum required levels for each modulation, taking into account a noise figure F = 10 dB and an implementation margin of 5 dB. However, modern devices improve these values and some receivers can successfully detect signals at even -100 dBm using 1 Mbps [26] (standard sets this limit at ≤-80 dBm).

Table 3.1 shows typical values for the receiver sensitivity and the distance at which these P_r values are obtained (with different propagation models), with a confidence of 95% (5 dB margin above the required minimum). In other words, Table 3.1 presents the cell radius of a WLAN, depending on the modulation and the propagation model.

3.1.4 Multipath channel

Different radio propagation phenomena result in a radio signal reaching the receiver by two or more different paths. This is known as *multipath propagation*. The phenomena causing multipath propagation is diverse: atmospheric ducting, ionospheric reflection and refraction. In the scenarios where WLANs are found, the presence of obstacles, the mobility of users and the possible lack of a line of sight between transmitter and receivers causes the transmitted waves to interact with surface irregularities via diffraction, scattering, reflection and absorption.

Radio waves reflect on the obstacles they meet (walls, furniture, floors, etc.). The echoing of a transmitted signal off those obstacles results in a series of signals reaching the receiver's antenna at different times, with different amplitudes and phases due to the different paths traversed by the different reflections. This produces constructive and destructive interference. The resulting signal combination fades rapidly. These rapid random fades can be modeled using different distributions, depending on the environment [295]. Rayleigh's distribution is used in heavily built-up urban environments with no dominant propagation along a line of sight between the transmitter and receiver. On the other hand, if there is line of sight and therefore one of the paths is much stronger than the others, Rician fading may be more applicable.

The mathematical models used to describe this behavior can be presented by the impulse response. Assume the transmission of an ideal Dirac pulse at time $t = 0$: $x(t) = \delta(t)$. Multiple pulses will be received at the receiver due to the presence of multiple paths. As explained above, each pulse will arrive at a different time, with a different amplitude. The received signal can be expressed as:

$$h(t) = \sum_{l=0}^{L-1} h_l \delta(t - \tau_l) \tag{3.9}$$

where h_l and τ_l are the l_{th} path complex gain and time delay, respectively. How h_l and τ_l are obtained differs from one model to another. The Naftali channel model is consistent in 802.11 systems [97]. With the use of this model, the channel impulse response is composed of complex samples using random uniformly distributed phase and Rayleigh distributed magnitude.

The time difference between the arrival moment of the first multipath component (typically the Line of Sight component) and the last one is called "delay-time spread" ($T_{DS} = \tau_{L-1} - \tau_0$). More precisely, the maximum delay-time spread is the total time interval during which reflections with significant energy arrive (the last impulse is the first one that allows a particular amount of the total transmitted power to be received, e.g. 99%). The term *Delay Spread* is then defined as the standard deviation value of the delay of reflections, weighted proportional to the energy in the reflected waves.
Delay spread introduces a phenomenon known as intersymbol interference (ISI) to the

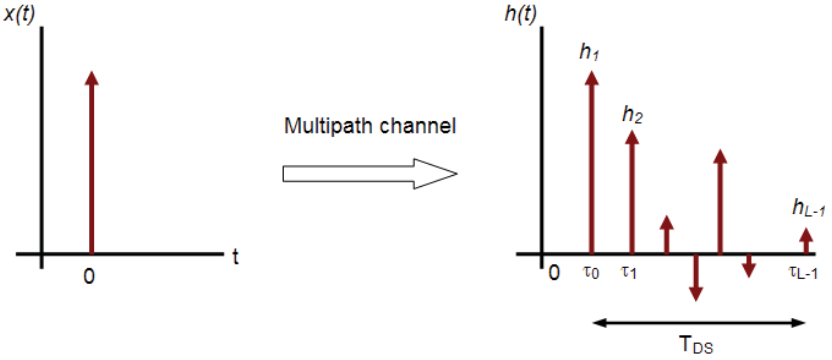

Figure 3.2: Impulse response of a multipath channel

receiver. To unravel the signal, the delay spread must be less than the symbol rate. Otherwise, a portion of the delayed signal will spread into the next symbol's transmission. Doing so requires a base band processor implementing the functions of an equalizer.
The ISI produces errors in reception. For this reason, the amount of delay spread that a system can tolerate is limited and depends on the modulation; the higher the symbol rate, the lower the admitted delay spread. The typical supported delay spread values for BER $\leq 10^{-5}$ range from 50 to 500 ns (for 1 Mbps). Note that a time difference of 50 ns corresponds to a path length difference of 15 m. In macro-cellular mobile radio, delay spreads are mostly in the range from less than 0.2 μs in open areas, 0.5 μs in suburban areas, to 3 μs in urban areas. In indoor channels, the delay spread is usually smaller (shorter paths), and rarely exceeds 300 ns.

In order to reduce the multipath effect, directive antennas can be used, at the cost of losing coverage. Circular wave polarization antennas (helicoidal antenna) are able to cancel the first reflections quite well. As mentioned above, a good equalizer is needed to reduce the harmful effect of ISI. However, we could benefit from multipath propagation if we use constructive interference to combat fading. Antenna diversity along with RAKE receivers are an efficient way of doing so [50, 205]. A RAKE receiver extracts multipath replicas with several correlators; then, a maximum ratio combining is applied using delay and weighting summation. In other words, the RAKE receiver uses several "sub-receivers", each one independently decoding a single multipath component; at a later stage the contribution of all "sub-receivers" are combined in order to make the most use of the different multipath components.

3.1.5 Measurement campaigns

In order to characterize path loss in a real scenario, different measurements were carried out during the development of the work presented in [207], supervised by the author of this thesis.

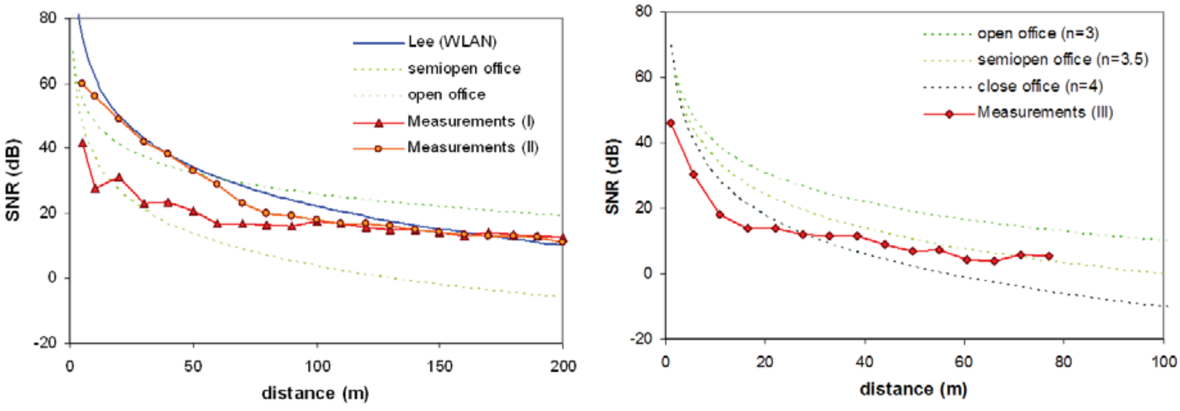

(a) Outdoor measurements

(b) Indoor measurements

Figure 3.3: Measured SNR in different outdoor scenarios

The first scenario consisted of an open space located on the beach of Castelldefels. One AP was placed 1 m from the ground in a fixed location. A laptop computer, associated to the AP, was used to measure the SNR as it moved away from the AP. A new measurement was taken every 10 m. The second set of measurements was taken in an urban/suburban area; an avenue surrounded by low buildings (2 - 4 floors). In both cases the transmitted power P_t was 15 dBm and the noise measured at the receiver -95 dBm. In addition, similar measurements were taken in an indoor scenario which could be defined as a semi-open office (90x5 m corridor, concrete walls, no doors and sparsely furnished). Figure 3.3 compares those measurements with the propagation models studied in previous subsections.

Surprisingly, the open environment presented more absorption than the urban scenario, especially throughout the first tens of meters. This is probably due to the high humidity and the proximity of the seashore in the first scenario. Moreover, in a line of sight scenario over a flat terrain such as the first one, we can consider only two paths for the signal: direct path and a ground reflection path. The contribution of these two signals at the receiver's antenna produce oscillations in the received power caused by the constructive and destructive combination of the two rays. For this reason, the relationship between SNR and distance depicted in Fig. 3.3(a) presents such variations in the short distances. The figure also shows that the propagation losses measured in the urban environment fit in with those predicted by the Lee (WLAN) model [149]. It is also worth noting that both SNR measurements tend to converge as the distance is increased. On the other hand, indoor measurements showed that semi-open or closed office models are well suited with exponents between 3 and 4.5. According to the measurements obtained in the indoor scenario described above, a dynamic selection of the exponent (e.g. $20 < d < 40m \rightarrow \gamma = 4; 40 < d < 60m \rightarrow \gamma = 3.5$, etc.) would provide more accurate results.

3.2 Error models

The probability that a symbol is erroneously received in the presence of Additive White Gaussian Noise (AWGN) is given by the following expression:

$$P_s = Q\left(\sqrt{\frac{d^2}{2N_0}}\right) \tag{3.10}$$

where d is the minimum euclidean distance between any two points in the constellation diagram of the modulation. N_0 is the noise spectral density (in W/Hz). The $Q(x)$ function is a convenient way to express right-tail probabilities for normal (Gaussian) random variables. Given that $x \in \Re$, $Q(x) \in \{0, 1\}$ is defined as the probability that a standard (zero mean and unit variance) normal random variable exceeds x.

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\left(t^2/2\right)} dt \tag{3.11}$$

For an implementation of the Q-function in a system model, a n_{th} order approximation can be used to avoid an infinite integral (c.f. [75, 234]).

3.2.1 OFDM

The IEEE 802.11a and .11g define a different set of modulations for the PHY layer (see Table 2.2). BPSK is used for 6 and 9 Mbps rates, QPSK for 12 and 18 Mbps and different flavors of QAM for the rest of the PHY modes (from 24 to 54 Mbps).

In BPSK, the phase of a constant amplitude carrier signal is switched between two values separated by 180$^{\text{o}}$, corresponding to binary 1 and 0. The two signal points of a BPSK are placed at a distance $\sqrt{E_b}$ from the origin (E_b is the energy per bit) in the constellation. The distance d between adjacent points is thus $2\sqrt{E_b}$. Substituting this into eq. 3.10, and taking into account that in BPSK one symbol corresponds to one bit, the probability of bit error in an AWGN channel is computed as:

$$P_b^{BPSK} = Q\left(\sqrt{\frac{2E_b}{N_0}}\right) \tag{3.12}$$

In M-ary PSK modulations, the amplitude of the transmitted signal is constant. By combining phase and amplitude variations, more constellation points can be drawn; in other words, more bits per symbol are admitted while the distance between adjacent points is not drastically reduced. These are the principles of QAM. As a matter of fact, PSK modulations could be considered as a special case of QAM, in which the magnitude of the modulating signal is kept constant. A brief consideration reveals that 2-QAM and

4-QAM are in fact BPSK and QPSK, respectively. However, this variation of amplitude entails that M-ary QAM does not have a constant energy per symbol, nor does it have a constant distance between symbols. For this reason, the error probability formulation can be provided according to the minimum energy per symbol [155] or in terms of the average energy per symbol [300]. Following [248], the equations for the symbol error probability for an M-ary QAM with an average energy per symbol E_s and an even number of bits per symbol (e.g. M = 4, 16, 64), are most easily expressed in a *per carrier* sense:

$$P_{\sqrt{M}-sym} = 2 \quad \left(1 - \frac{1}{\sqrt{M}}\right) \quad Q\left(\sqrt{\frac{3}{M-1}\frac{E_s}{N_0}}\right) \tag{3.13}$$

where $P_{\sqrt{M}-sym}$ is the symbol error probability for the $\sqrt{M}$-ary Pulse Amplitude Modulation (PAM). Then:

$$P_{sym}^{QAM}(M) = 1 - \left(1 - P_{\sqrt{M}-sym}\right)^2 \tag{3.14}$$

The BER will depend on the exact assignment of bits to symbols, but if gray coding is used, with equal bits per carrier and $k = log_2(M)$ representing the number of bits per symbol:

$$P_b^{QAM}(M) = 1 - (1 - P_{\sqrt{M}-sym}(M)/k)^2 \tag{3.15}$$

However, as in [252], we consider a binary convolutional coding with a hard-decision Viterbi decoding, and we also assume that the errors at the channel input are independent. Under these assumptions, the upper bound for the Packet Error Ratio (PER) was given in [250]:

$$PER_m(l, SINR) \leq 1 - (1 - P_u^m(SINR))^l \tag{3.16}$$

The PER depends on the SINR per symbol (E_s/N_0) and also on the length of the packet (l bits) for any PHY mode M used for its transmission. The term P_u^M is the union bound on the first-event error probability. The Viterbi algorithm decodes by comparing the sequence of received symbols to the symbol sequence associated with each path and choosing the most likely path (i.e., the one whose sequence is closest in Hamming distance to the received sequence). A first-event error is said to occur at node i if the decoder chooses a path that first deviates from the top line of the trellis at node i. Roughly, we can say that P_u^M is the BER that includes the effects of the coding gain applied to PHY mode M (see Table 2.2). Then, P_u^M is given by:

$$P_u^m(SINR) = \sum_{d=d_{free}}^{\infty} a_d P_d \tag{3.17}$$

where d_{free} is the free distance of the convolutional code selected in PHY mode M, a_d is the total number of error events of weight d, and P_d is the probability that an incorrect

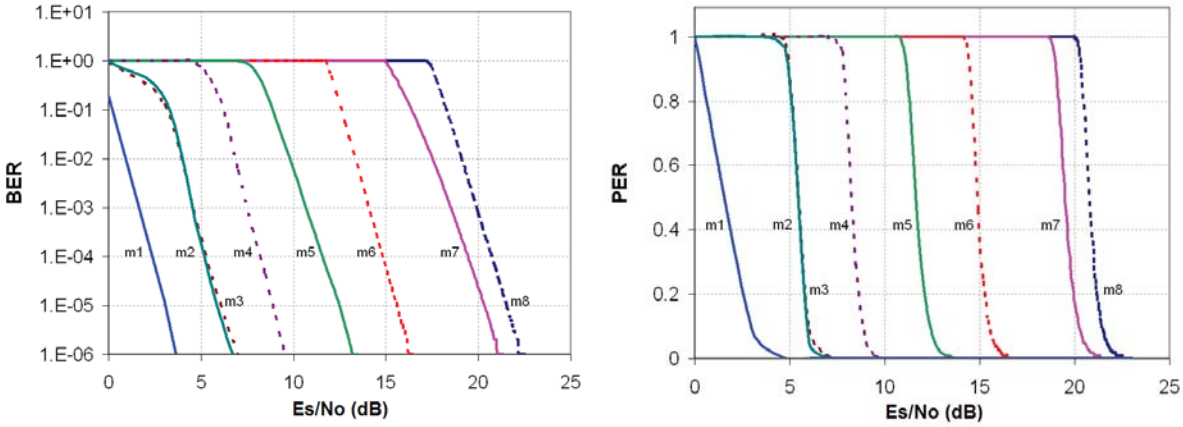

(a) BER (P_u^m) vs SINR for OFDM modes

(b) PER vs SINR for 1024-Byte packets

Figure 3.4: Error model for the different OFDM PHY modes

path at distance d from the correct path is being chosen by the Viterbi decoder. Both d_{free} and a_d depend on the code rate of the PHY mode M and are given in [151]. Finally, when the hard-decision decoding is applied, P_d is computed as a function of P_b as follows:

$$P_d = \begin{cases} \sum_{k=(d+1)/2}^{d} \begin{pmatrix} d \\ k \end{pmatrix} P_b^k (1-P_b)^{d-k} & \text{if } d \text{ is odd} \\ \frac{1}{2} \begin{pmatrix} d \\ d/2 \end{pmatrix} P_b^{d/2} (1-P_b)^{d/2} + \sum_{k=d/2+1}^{d} \begin{pmatrix} d \\ k \end{pmatrix} P_b^k (1-P_b)^{d-k} & \text{if } d \text{ is even} \end{cases} \tag{3.18}$$

P_b depends on the PHY mode selected and is computed according to eq. 3.12 or eq. 3.15. Of course P_b depends on the SINR. The resulting BER for the different IEEE 802.11a/g were tabulated for different E_s/N_0 values in [208]. These values were used to plot the BER vs. SINR Figure 3.4(a). Figure 3.4(b) shows the PER according to eq. 3.16 and taking $l = 8192$ bits (1024 Bytes).

3.2.2 DSSS

As introduced in 2.4.1.2, the modulation scheme for the IEEE 802.11 DSSS at 1 Mbps is DBPSK and DQPSK for the 2 Mbps rate. Differential modulation are slightly worse in terms of BER if compared with their non-differential counterparts. On the other hand, differential PSK avoids the need for possibly complex carrier-recovery schemes to provide an accurate phase estimate. It is commonly accepted that with differential modulation, the noise power at the receiver is doubled [256], though it is a pessimistic approach, specially for high SINR values. Thus:

$$P_{DPSK-chip} = Q\left(\sqrt{\frac{d^2}{4N_c}}\right) \tag{3.19}$$

where N_c is the noise energy per chip. Plotting the signal constellation of a BPSK modulation, the signal points are placed at a distance $\sqrt{E_c}$ from the origin (E_c is the energy per chip). Thus, the distance d is $2\sqrt{E_c}$ for DBPSK, and $\sqrt{2E_c}$ for DQPSK. This leads to a chip error probability of:

$$\begin{aligned} P_{DBPSK-chip} &= Q\left(\sqrt{E_c/N_c}\right) \\ P_{DQPSK-chip} &= Q\left(\sqrt{E_c/2N_c}\right) \end{aligned} \tag{3.20}$$

To include the effect produced by the processing gain of DSSS, the squared distance is summed over each chip. In the case of IEEE 802.11 DSSS modulations, the 11-chip spreading word results in the squared distance being multiplied by 11 and therefore:

$$\begin{aligned} P_{sym}^{DBPSK} &= Q\left(\sqrt{11.E_c/N_c}\right) \\ P_{sym}^{DQPSK} &= Q\left(\sqrt{5.5.E_c/N_c}\right) \end{aligned} \tag{3.21}$$

This is the Symbol Error Ratio (SER) of the DSSS modulations with 1 and 2 Mbps. For the 1 Mbps case, BER = SER, since each symbol encodes one single bit. On the other hand, for the 2 Mbps case, each symbol encodes two data bits. Nonetheless, given that Gray coding is used, a decoding error between adjacent DQPSK constellation points (more probable than errors between opposing points), produces only 1 bit error, i.e. one symbol error corresponds to one bit error, so again we have that in practice, BER = SER.

This analysis corresponds to a theoretical AWGN channel. Different formulations can be found in the literature that take into account the effects of a more complex channel model (e.g. see [212]). Furthermore, actual implemented receivers may use coherent demodulation to improve BER. For example, the Intersil's HFA3861B [9] baseband processor uses differential demodulation for the initial acquisition portion of the message processing, and then switches to coherent demodulation for the MPDU demodulation. According to the analysis in [270], the improvement for DBPSK is near 1 dB compared to non-coherent demodulation.

3.2.3 CCK

According to [272], the error performance of a CCK signal can be accurately modeled following an M-ary Quaternary Orthogonal Keying Modulation (MQOK). However, the BER of a CCK can be determined by treating the modulation as a block code, and hence a less complex formulation can be obtained. Again, in the presence of an AWGN channel, the symbol error rate is obtained through eq. 3.10. The general expression for SER is:

$$P_{sym}^{CCK} \approx \sum Q\left(\sqrt{2 \quad E_b/N_0 \quad R_c \quad W_m}\right) \tag{3.22}$$

where R_c is the code rate. In 11 Mbps CCK modulation, 8 chips encode 8 data bits and then $R_c = 1$. For 5.5 Mbps, 4 data bits are encoded in 8 chips, thus $R_c = 1/2$. W_m is the codeword distance. The summation must be done over all (8-chip length) 256 codewords (for 11 Mbps), or over all (8-chip length) 16 codewords (5.5 Mbps case). If each symbol encodes k data bits (4 for 5.5 Mbps and 8 for 11 Mbps), the relationship between BER and SER can be expressed as follows:

$$P_b^{CCK} \leq P_{sym}^{CCK} \frac{2^{k-1}}{2^k - 1} \tag{3.23}$$

It can be shown that the minimum squared distance for CCK (given the codewords obtained following eq. 2.1) is $d^2 = 8E_b$ [192]. For each symbol, there is an average of 24 neighbors at this distance, when the 256 11 Mbps codewords are considered. Likewise, one can compute neighboring distances for remaining codewords in order finally to obtain the performance expression for 11 Mbps (3.24) and 5.5 Mbps (3.25):

$$\begin{aligned} P_{sym}^{CCK11} \leq & 24Q\left(\sqrt{4E_b/N_0}\right) + 16Q\left(\sqrt{6E_b/N_0}\right) + 174Q\left(\sqrt{8E_b/N_0}\right) + \\ & + 16Q\left(\sqrt{10E_b/N_0}\right) + 24Q\left(\sqrt{12E_b/N_0}\right) + Q\left(\sqrt{16E_b/N_0}\right) \end{aligned} \tag{3.24}$$

$$P_{sym}^{CCK55} \leq 14Q\left(\sqrt{8E_b/N_0}\right) + Q\left(\sqrt{16E_b/N_0}\right) \tag{3.25}$$

It is necessary to point out that these formulas are not accurate for small values of the SINR [12], but CCK modulations are seldom used with SINR $<$ 7 dB (c.f Section 4.3.3). In [192], the effects of quantization are added to increase the accuracy of the above formulation. In [205], a different approach is given to analyze CCK performance under multipath fading channel.

Figure 3.2.3 shows the performance of the different IEEE 802.11b modes. It is interesting to note from 3.5(a) that CCK 5.5 Mbps is more energy-efficient than 2 Mbps DSSS since for a given BER, the CCK modulation requires a lower E_b/N_0. This fact does not entail that 5.5 Mbps is preferable for any SINR. Recall that the E_b/N_0, measured at the receiver input terminals, represents the SINR *per bit*. However, the energy per symbol required to guarantee a given BER is always lower for the 2 Mbps mode, since the latter contains less bits per symbol. This difference is shown in Fig. 3.5(b): the PER vs E_s/N_0 is better for 2 Mbps.

This performance improvement (in terms of E_b/N_0) is due to the embedded coding properties of the spreading modulation used for 11 and 5.5 Mbps. As explained in [50, 51],

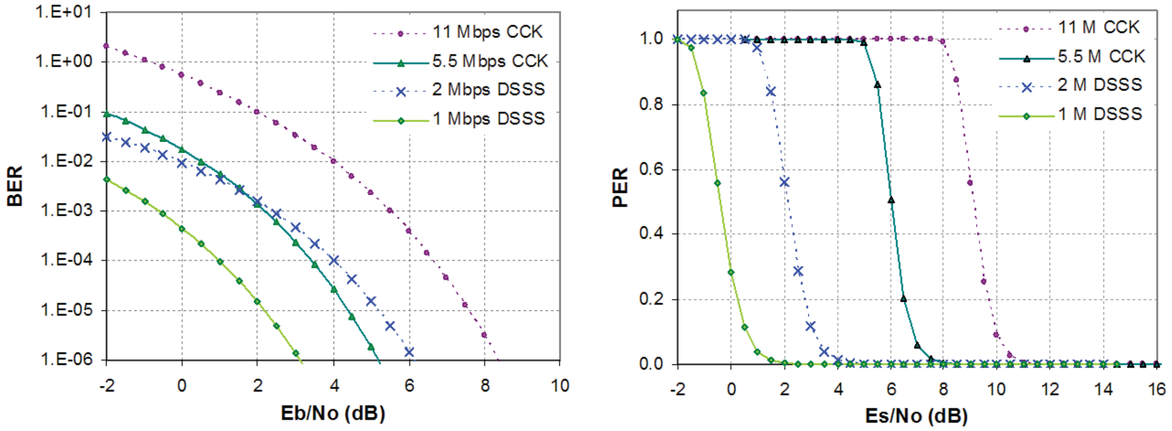

(a) BER vs SINR for IEEE 802.11b modes

(b) PER vs SINR for 1024-Byte packets

Figure 3.5: Error model for the different DSSS/CCK PHY modes

the modulation basically ties several bits together so that the receiver makes a symbol decision: if a symbol is in error, then all the bits in that symbol are suspect, but not all will necessarily be erroneous. In other words, the energy required to make all the bit decisions of a symbol independently and correctly is higher than the energy required to make a symbol decision correctly.

3.3 Interference and coexistence

The study presented in [45] states that in most cases the density of WLAN nodes in an area is such that administrators are not able to ensure an innocuous coexistence (many interfering sources and a limited number of non-overlapping frequency channels). It is clear that, with an increasing number of neighboring nodes, the undesired effect of interference becomes more problematic, affecting the network performance. Although intermodulation and intersymbol interference constitute an important concern for wireless radio system designers, network administrators usually focus on two types of interference [246]: *co-channel interference*, which is caused by undesired transmissions carried out on the same frequency channel; and *adjacent channel interference*, produced by transmissions on adjacent or partially overlapped channels.

Furthermore, due to the use of ISM bands, the performance of WLANs can be hampered by other (non-802.11) sources of interference, such as cordless phones, microwave ovens, pulse radar signals, Bluetooth or Zigbee devices, etc. Proposed solutions intended to reduce this kind of interference are commonly based on new PHY and MAC layer mechanisms (e.g. adaptive filters, frequency hopping schemes, etc.). The next subsection serves as a review of these techniques. On the other hand, the interference produced by neighboring 802.11 cells can be effectively alleviated with an intelligent radio resource management, which is the topic studied in this thesis. Therefore, in this section we will focus on the study and characterization of interference, in which the RF sources are 802.11 compliant.

3.3.1 Non-802.11 interference

Interference from microwave ovens have been widely described in the literature [174, 228, 229]. According to National Telecommunications and Information Administration (NTIA) reports [143, 144], residential ovens radiate 16 to 33 dBm with a center frequency around 2.45 and 2.46 GHz, and occupy a band about 20 MHz wide. Therefore, the best way of avoiding such interference is to limit the available channels to numbers 1 to 6. However, radiation of a microwave oven is cyclic: typically, they are active for a time period of about 8 ms over a power cycle of 20 ms when the power supply frequency is equal to 50 Hz, or of 16 ms when the power supply frequency is equal to 60 Hz. In this way, several overlap avoidance (OLA) techniques (usually based on simple traffic scheduling) could be applied to allow coexistence even in the bands affected by the microwave radiation (e.g. see [99, 265]).

Interference caused by Bluetooth systems and other Wireless Personal Area Network (WPAN) technologies using ISM bands has also received much attention. Coexistence between these systems and WLANs has been studied by the IEEE 802.15 TG2. Their recommended practices were published in [12]. According to this standard, coexistence can be promoted either via collaborative or non-collaborative mechanisms. Collaborative mechanisms imply that there is a common point where the WLAN and the WPAN meet in order to exchange information related to their traffic transmissions. That is to say, collaborative mechanisms are intended to be used when at least one WLAN station and one WPAN device are placed (co-located) within the same physical unit. When co-located, there needs to be a communication link between the WLAN and WPAN devices within this physical unit. On the other hand, if such a communication link does not exist, non-collaborative mechanisms could be used. Independently of the IEEE 802.15.2 work, other mechanisms were proposed that can be deployed either in a collaborative or in a non-collaborative scenario [99].

The IEEE 802.15 TG2 discussed two collaborative proposals at MAC level. Packet Traffic Arbitration (PTA) involves the use of a centralized controller that monitors Bluetooth and the 802.11 traffic; the controller provides per-packet authorization of all transmissions and uses its knowledge of the 802.11 and Bluetooth activity to predict collisions. When a collision is likely to occur, PTA schedules transmissions based on simple rules determined by the packet types (e.g. synchronous connection oriented (SCO) Bluetooth traffic has higher priority than any 802.11 asynchronous connectionless (ACL) data packets). With Alternating Wireless Medium Access (AWMA), the 802.11b beacon-to-beacon interval is subdivided into two subintervals: one subinterval for 802.11 and the other subinterval for Bluetooth. Since each radio has its own subinterval, both radios will operate properly, due to total orthogonality. The two MAC layer techniques can be integrated with a collaborative PHY solution, the deterministic interference suppression. The key idea is that because the Bluetooth signal can be considered a narrowband interferer for the 802.11 DSSS signal, we can put a null in the 802.11 receiver at the frequency of the Bluetooth signal. Logically, the hopping sequence and its timing must be known by the 802.11 receiver.

If such collaboration is not possible, three different alternatives are present in [12]: adaptive interference suppression of IEEE 802.11 devices, adaptive packet selection, and packet scheduling for ACL links. Two other non collaborative mechanisms are included as informational (packet scheduling for SCO links and adaptive frequency-hopping (AFH) for the IEEE 802.15 devices).
IEEE 802.15.1 specifies a variety of packet types with different combinations of payload length, slots, degree of error protection, etc. The adaptive packet selection and the packet scheduling mechanisms take advantage of this characteristic. By selecting the best packet type according to the channel condition of the upcoming frequency hop, better data throughput and network performance may be obtained. In addition, packet transmission could be carefully scheduled so that the IEEE 802.15 devices transmit during hops that are outside the WLAN frequencies and refrain from transmitting while in-band.

The collaborative interference suppression method previously described requires an IEEE 802.15 receiver co-located with the WLAN receiver. If this collaboration is not enabled, a new approach based solely on signal processing in the WLAN's PHY is recommended. In this method, the WLAN has no a priori knowledge of the timing or frequency used by the WPAN system, and it uses an adaptive filter to estimate and cancel the interfering signal.
Finally, according to the adaptive frequency hopping scheme, Bluetooth channels are classified as "good" or "bad", depending on RSSI, packet loss ratio, etc. "good" and "bad" channels are then used intelligently to reduce the probability of overlap in frequency with the 802.11 signal, maximizing the use of "good" channels.

Nevertheless, and as stated before, there are many potential interference sources of a different nature. Causes and effects of non-802.11 interference are investigated in [150] from a more general point of view. According to the conclusions derived in this work, the best strategy that is intended to minimize this kind of RF interference (including malicious 802.11-based *jammers*) is achieved by implementing some sort of channel hopping [59,233]. That is, a new mechanism is required that allows the members of a WLAN to perform periodical frequency hops synchron

3.3.2 Receiver design

If we want to characterize and analyze the causes and effects of interference, we first have to understand the operation of a receiver. Some of their characteristics have been introduced before: Section 3.1.3 presents the effect of noise on receiver sensitivity, and 3.1.4 explains the consequences of delay spread caused by multipath channels.

This subsection is not intended to settle a discussion on different receiver architectures (e.g. super-het vs. Direct Conversion), but rather to present some common limitations that affect any receiver's performance. According to [150], these limitations can be classified in three different categories: timing recovery, dynamic range selection and PLCP header processing.

The reception process can be explained in a sequence, following the PLCP headers shown in Fig. 2.5. First of all, the reception of a specific SYNC bit-pattern triggers the energy detection functions that alert the receiver of an incoming transmission. This preamble is also used to obtain symbol timing. The SFD bit sequence indicates the start of the PLCP header (see 2.4.1.1 and 2.4.2.2). The PLCP header is then protected by a CRC. Using this CRC, the receiver generates a physical layer error if it detects a corrupted header. Following the PLCP header is the MAC frame, which includes a new CRC that protects the MAC contents (see Fig. 2.2). If the MAC is corrupted, a separate error is generated.

All these PLCP and MAC processes are affected by the presence of interference. The next subsections provide explanations for these effects.

3.3.2.1 Timing recovery and header processing

After amplification and filtering of the RF signal, a down conversion mixer provides baseband signal to an Analog-to-Digital Converter (ADC). The ADC transforms the analog signal into bit samples, and these samples are processed to recover the sender's clock so that the following receiver components can work on aligned signal samples. Time synchronization begins when the receiver detects the SYNC pattern of a DSSS/CCK (128 bits for long preamble, 56 for short preamble), or the first 10 symbols of an OFDM signal. If the receiver fails to lock onto the sender's clock, the energy detected will be reported as a busy channel, but the receiver will not be able to recognize and decode a valid modulated frame.

As mentioned before, the PLCP header contains information required to successfully decode the frame. For example, the *LENGTH* field provides the duration of the MPDU and the *SIGNAL* field conveys the modulation and coding rate applied to the remaining bits of the frame. In consequence, if the PLCP header is corrupted, a PHY header checksum error is generated and the whole frame is discarded.

3.3.2.2 Dynamic range selection

The received signal strength falls over a large range of power levels. Typically, the weakest received signal falls below -70 dBm, and the strongest signal may be above -20 dBm. This

represents a range of more than 50 dB; that is, received signal strength can increase or decrease by a factor of 10^5. In order to support this wide variety of input levels, the signal is internally normalized to the fixed range admitted by the ADC unit.
In the receiver chain, there is an Automatic Gain Control (AGC) unit that samples input voltage levels during the PLCP preamble processing, and controls the gain of RF amplifiers. The gain of these low noise amplifiers is adjusted so that the entire ADC input range is occupied. The ADC is then able to identify up to 64 voltage levels, converted to 6-bit words (typically, e.g. [9]).

The input level of most 802.11 amplifiers saturates at around -20 dBm, even though a well designed amplifier can still operate up to an input level of -10 dBm [15]. Some systems are also able to bypass the amplifier when the input signal exceeds a given threshold (e.g. -10 dBm). This allow input signals up to 4 dBm at the cost of a reduced receiver sensitivity. In the presence of input signals above the saturation range, the amplifier will cease providing gain. Instead, it will introduce non-linear distortion into the signal. Therefore, even a narrow in-band interference will be able to outwit the spreading spectrum protection if it breaks the saturation point of the low noise amplifiers.

The above processes are done once per packet during the PLCP preamble reception. This means that if the interference or the noise floor varies substantially after the gain control is performed, the voltage levels at the ADC input are not adjusted. In consequence, even a low-powered and narrow in-band interference is able to subvert the correct reception of a frame if it is received at different levels during and after the frame's preamble. When such an interference is received after the gain control has attenuated a strong signal, it can cause the ADC to overflow. Similarly, if that interference is removed after the preamble, the output of the ADC underflows [150]. Hence, the effects of such interference is the increase of packet loss due to PHY layer CRC errors, or, if the PLCP preamble and header were not affected by the interference, in the data payload.

3.3.3 Channel scheme

Interference issues in IEEE 802.11 WLANs are closely related to the definition of frequency channels in the different regulatory domains. The spectrum available for ISM communications is scarce and provides insufficient room for a small number of simultaneous WLAN transmissions.
After showing the frequency spectrum devoted to ISM communications in different countries, this section shows the appearance of an IEEE 802.11 signal in the frequency domain.

3.3.3.1 Regulatory domains

All WLAN (with the exception of infrared-based networks) deployments that make use of unlicensed spectrum fall into three basic ISM frequency bands: 900 MHz, 2.4 GHz, and 5 GHz. The different regulatory domains define which frequencies and channels may be used. These definitions are in an ever-changing state.

The 900-MHz band was the first area for which spread-spectrum WLANs were developed. The availability of inexpensive, small RF components developed for use in the cellular phone industry helped the early deployment of those WLANs. However, it was confined to International Telecommunication Union (ITU) Region 1 (mainly the Americas) and had a very limited bandwidth. Due to the limited span that was available, running the higher data rate (2 Mbps) reduces the number of channels to one that incorporates the entire band. These two major drawbacks made the IEEE look at the 2.4 and 5 GHz ISM bands on the ITU Radio Regulations [31], for the development of the 802.11 standards [49].

The 2.4-GHz band was generally available in almost every major country worldwide. It initially provided for data rates of only up to 2 Mbps, but soon the 802.11b and 11g amendments increased those rates. Because the frequency scheme is identical between the initial 2.4-GHz 802.11, the 802.11b, and the 802.11g specifications, most countries that permitted operation for the early 2.4-GHz 802.11 devices also permitted the 802.11b and 802.11g products.
The center frequency of the first channel is defined at 2.412 GHz and, once spread, the resulting signal occupies a bandwidth of about 22 MHz (see Section 3.3.3.2). In addition, the available channels are defined with 5 MHz separation between consecutive carriers, giving rise to the need to use at least five channels of separation to guarantee that two simultaneous transmissions do not interfere with each other. In the North American domain (NA), there was a need to limit the upper channels because of a very tight restriction for RF signals falling outside the band. Therefore, there were only 11 channels specified (upper channel centered at 2.462 GHz).
In the European Telecommunications Standards Institute (ETSI) domain, 2.4 GHz devices must conform with EN 300 328 standard [23]. According to this standard, the permitted band ranges from 2.4 to 2.4835 GHz. Consequently the lower 11 channels are identical to the NA scheme, but also two additional channels are allowed (at 2.467 and 2.472 GHz).
In Japan, a very strict regulation limited WLAN usage in the 2.4 GHz band to only one channel (at 2.484 GHz), and that channel was incompatible with any of the ETSI or NA channel definitions. Some years later, the Japan TELEC changed the regulations, permitting operation of the 13 ETSI channels plus the old single Japan channel, thus providing for 14 channels [49]. Consequently, whereas there are up to 19 non-overlapping channels in the 5 GHz band, in the 2.4 GHz band only three out of 14 (11 in USA) are non-overlapping (traditionally, channels 1, 6 and 11), as shown in Fig. 3.6.

While in Europe, the 5-GHz band was initially used for the ETSI HiperLAN specification, the competing standard IEEE 802.11a, which was released in 1999, got hold of the market. The 802.11a specification defined several different channel groups within the 5 GHz band.

As previously mentioned, the 5-GHz band is broken down into several different channel

groups. In the United States, these are referred to as the Unlicensed National Information Infrastructure (UNII) bands. The three bands or groups UNII1, UNII2, and UNII3 permit operation in the 5.215 to 5.225 GHz, 5.225 to 5.235 GHz, and 5.725 to 5.825 GHz frequency ranges, respectively. After the recent changes in regulations, a new frequency band is now available ranging from 5.470 to 5.725 GHz.
In Europe, the ECC/DEC/(04)08 [16] establishes two sub-bands: 5.150 to 5.350 GHz and 5.470 to 5.725 GHz. Devices in these bands must conform to the EN 301 893 standard [27]. According to this standard, the first band is only permitted for indoor communications with a mean EIRP of 200 mW and DFS together with TPC (see 7.1.1) are additionally required above 5.250 GHz. The second band (5.470 – 5.725 GHz) can be used for both indoor or outdoor deployments if DFS and TPC are implemented and the mean EIRP is ≤ 1 W.

3.3.3.2 Spread signals

IEEE 802.11 networks operate in the 5 GHz (.11a) and 2.4 GHz (.11b/g) unlicensed frequency bands. Communications in these bands need to implement spread spectrum techniques and limit their transmitted power in order to minimize the impact of interference with other devices. The IEEE 802.11 defines different spreading techniques, but as mentioned earlier, the devices that can be found today on the market are based on DSSS and OFDM. In DSSS, the data at the sending station is combined with a higher-rate bit sequence that spreads the user data in frequency by a factor equal to the spreading ratio. The IEEE 802.11 standards specify the use of Barker codes (1 and 2 Mbps) and the use of CCK (5.5 and 11 Mbps) for the chip sequence in DSSS systems. The direct modulation effectively spreads the signal over a much wider bandwidth and its power spectrum can be described by the following equation:

$$P_{DSSS}(x) = \begin{cases} \left(\frac{sin(2\pi x)}{2\pi x}\right)^2 & \text{if } x \neq 0 \\ 1 & \text{if } x = 0 \end{cases} \tag{3.26}$$

where x is a function of the center frequency f_c, and the bandwidth of the main lobe b_m: $x(f_c, b_m) = (f - f_c)/b_m$. A general rule of thumb for DSSS systems is that the null

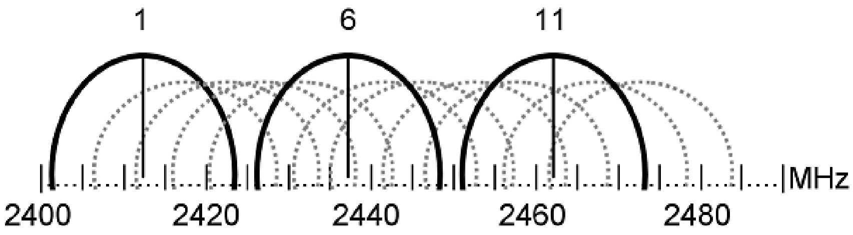

Figure 3.6: 2.4 GHz ISM band

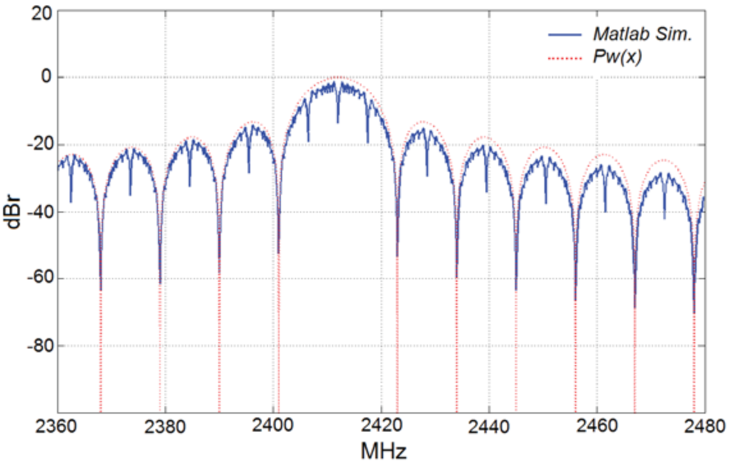

Figure 3.7: PSD of an unfiltered 802.11 DSSS signal

to null bandwidth is 2x the chip rate. Since all DSSS modulations use a chip rate of 11 Mcps, the null to null bandwidth of the spread signal (b_m) is 22 MHz. In Fig. 3.7, the spectrum defined by equation 3.26 is compared with Matlab simulations, consisting of a random symbol sequence that is created to represent the data to be transmitted; the data is then spread using the 11-bit Barker sequence, and finally, a carrier wave is applied.

The IEEE 802.11a/g standards [4] [13] specify an OFDM Physical Layer that splits an information signal across 52 separate sub-carriers. As seen in Section 2.4.2, recall that four of the sub-carriers are pilot sub-carriers that are used as a reference to disregard frequency or phase shifts. The remaining 48 sub-carriers provide separate wireless "pathways" for sending the information in a parallel fashion. The resulting sub-carrier frequency spacing is 0.3125 MHz (20 MHz/64) and the total bandwidth is 20 MHz, but only 16.6 MHz are actually occupied. According to the analysis in [204], the spectrum of the signal can be obtained by summing the power spectra of all individual sub-carriers ($P_{sk}(f)$). This power density is obtained directly from the Fourier transform of the time-window function defined by the standard:

$$W(f) = T_s \frac{sin(\pi T_s f)}{\pi T_s f} \frac{cos(\pi T_{tr} f)}{1 - 4T_{tr}^2 f^2} e^{-j\pi T_s f} \tag{3.27}$$

where T_s is the symbol duration (4 μs), and T_{tr} is the transition time, about 100 ns. The values for T_s and T_{tr} are the same for all modulations. Then, for the k_{th} sub-carrier:

$$P_{sk} = \frac{|W(f - 0.3125k)|^2}{T_s}; \quad k = \pm1; \pm2; ...; \pm52/2; \tag{3.28}$$

Note that the spectrum of any OFDM sub-carrier is only affected by the symbol shaping window and symbol rate.

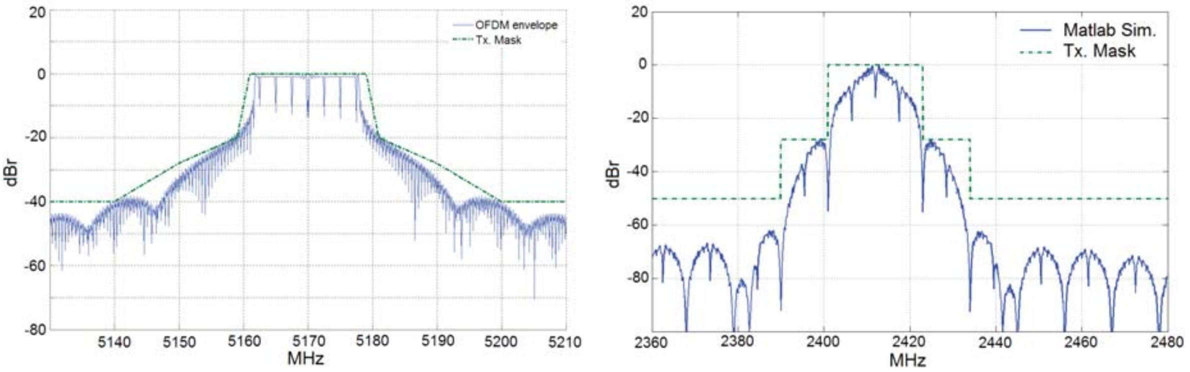

(a) Spectral Density of 802.11 OFDM signal

(b) Spectral Density of 802.11 filtered DSSS

Figure 3.8: IEEE 802.11 transmit spectrum mask

However, the IEEE 802.11 standard defines a transmit spectrum mask intended to limit the energy of the transmitted signal that invades adjacent channels: around f_c, the signal is unmodified; at frequencies beyond f_c± 11 MHz, the transmitted spectral products shall be less than -30 dBr, and -50 dBr for frequencies f_c± 22 MHz. For OFDM, the transmitted spectrum shall have a 0 dBr bandwidth not exceeding 18 MHz, -20 dBr at 11 MHz frequency offset, -28 dBr at 20 MHz frequency offset and -40 dBr at 30 MHz frequency offset and above (see Fig. 3.8). Consequently, a band-pass filter must be applied before transmitting to the medium. In Fig. 3.8(b), a Matlab simulation illustrates the effect of applying the spectrum mask by means of a 4th-order elliptic filter with 22 MHz of bandwidth and a stop band 50 dB down to a DSSS signal.

It is also interesting to observe the power spectrum of a real IEEE 802.11 transmission. For this reason, we include pictures from a spectrum analyzer[1] [20].
Figure 3.9 shows the spectrum measured for two IEEE 802.11 transmissions using channel 11 (centered at 2.462 GHz). The trace is obtained using the max-hold mode, that is, the picture shows the highest energy values measured during a two-minute transmission. The measured spectra corresponds to the signal at the receiver's antenna; in other words, it is the transmitted signal as it travels the wireless medium, once the transmission mask is applied, but before any filtering at the receiver. Figure 3.9(a) represents an IEEE 802.11g transmission using the 12 Mbps modulation. The presence of the multiple sub-carriers that characterize an OFDM signal can be clearly identified. Figure 3.9(b) is obtained after an IEE 802.11b transmission using 11 Mbps CCK.

As previously explained, one could draw the conclusion that the OFDM signal occupies less bandwidth than the DSSS/CCK mode, since the first uses 16.6 MHz, while the latter presents a null to null bandwidth of 22 MHz. However, the occupied bandwidth marker is commonly used to show where a specified percentage of the power lay on the trace, in this

[1] Rohde&Schwarz FSH6 handheld spectrum analyzer

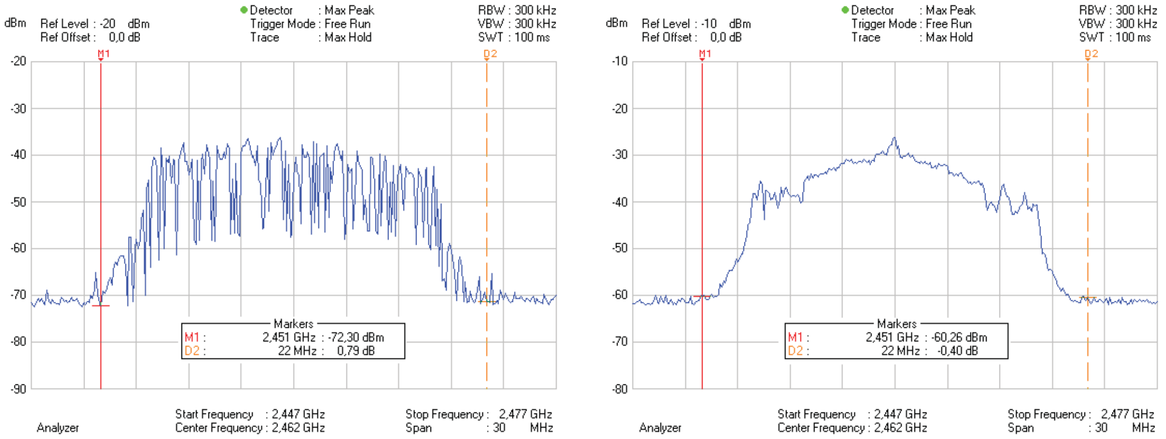

(a) Measured spectrum of an 802.11 OFDM signal (b) Measured spectrum of an 802.11 DSSS signal

Figure 3.9: Measured spectrum of different IEEE 802.11 transmitted signals

case 99 percent. Figure 3.10 shows the occupied bandwidth (area of the signal between light blue lines) measured using the spectrum analyzer. The analyzer reports that the 99% of the power transmitted by the OFDM device is spread over 16.1 GHz, while the DSSS/CCK transmission is concentrated in 14.7 GHz.

3.3.4 Adjacent channel interference

The presence of adjacent channel interference reduces the effective Signal to Interference and Noise Ratio (SINR), and therefore, the number of errors in reception is increased.

Similar to [225], we proposed in [129] a simple model to quantify the interference caused by transmissions in partially overlapped channels. The key idea of this model is to take an integral over the whole overlapping region of the interfering channels. In [225] this integral only involves the overlapping areas of the transmitter's and receiver's spectral masks, ignoring the real spectrum of the wave. However, in [129] we compute the power spectral density (PSD) of the filtered IEEE 802.11 signal in order to provide more accurate results.

Assume a receiver tuned to f_c, and a transmitter c channels apart that sends a signal represented by $P_w(f)$. Also assume that the two devices use an identical filter for both transmission and reception; if the filter's frequency response is represented by $F_{fc}(f)$, where f is the frequency in MHz, the overlapping energy of the interfering signal is computed as:

$$P_{int} = \int_{-\infty}^{+\infty} P_w(f) F_{fc}(f - 5c) F_{fc}(f) df \tag{3.29}$$

where $P_w(f) = P_{DSSS}(f_c, 22)$ for DSSS (see eq. 3.26) and $P_w(f) = \sum_k P_{sk}(f)$ for OFDM

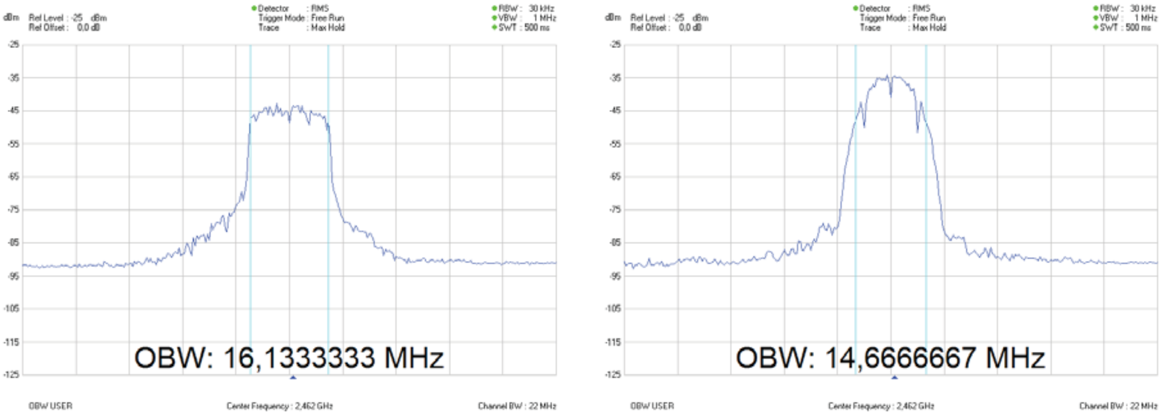

(a) Occupied bandwidth of an OFDM signal (99%) (b) Occupied bandwidth of a DSSS signal (99%)

Figure 3.10: Occupied bandwidth of different IEEE 802.11 signals

c **(ch. sep.)**	**0**	**1**	**2**	**3**	**4**	**5**
DSSS (Theor.)	0	0.28	2.19	8.24	25.50	49.87
DSSS (Sim.)	0	0.37	1.79	8.03	23.47	53.21
OFDM (Theor.)	0	0.55	2.46	6.60	34.97	51.87

Table 3.2: Attenuation values (in dB) for adjacent channels

signals (see eq.3.28). Note that for our purposes, the integral could be done only in the region of interest, i.e. $f_c\pm$ 22 MHz. Thus, we quantify the interference caused by 802.11 transmissions, according to the attenuation of the filter for a specific channel, normalized by the amount of energy the receiver would get if tuned to that channel. In other words, if we call P_0 the power that the receiver gets when both the receiver and the transmitter are on channel 1, and P_c is the new power obtained after moving the sender c channels away, the normalized loss factor is P_c/P_0.

Figure 3.11 shows graphically how adjacent OFDM (Fig. 3.11(a)) and DSSS (Fig. 3.11(b)) signals are filtered in reception. According to Fig. 3.11, P_c is the amount of energy received from channels 36 (OFDM) and 3 (DSSS). In other words, P_c is the area under green lines. Conversely, P_0 is the energy received from channel 34 (OFDM) and 1 (DSSS). P_0 is then the area under the blue lines.

Table 3.2 shows the attenuation values in dB ($P_0 - P_c$), obtained by means of Matlab simulations, and theoretically following equation 3.29 with perfect band-pass filters (bandwidth of 22 MHz and a stop band 50 dB down).

Note that simulated values for OFDM attenuation are not included, since the theoretical curves (eq. 3.27) were previously validated in [204]. The values in Table 3.2 are graphically represented in Fig. 3.12.

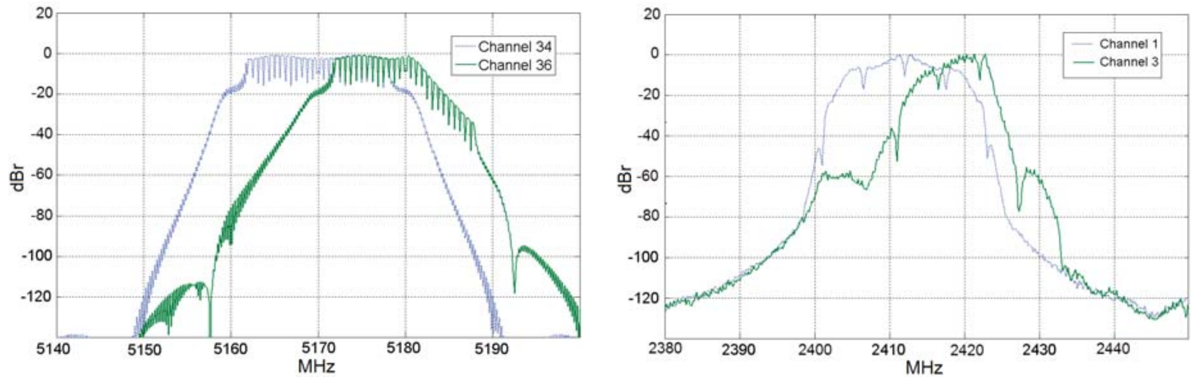

(a) OFDM Receiver tuned to ch. 34 filtering signals received in 34 and 36

(b) DSSS Receiver tuned to ch. 1 filtering signals received in 1 and 3

Figure 3.11: Filtering of adjacent IEEE 802.11 signals

For example, if a DSSS receiver is tuned to channel 1, it receives all transmissions in channel 1 ($c = 0$) without attenuation, but interfering transmissions on channel 4 are reduced by more than 8 dB ($c = 3$).

3.3.4.1 Effects of utilization on interference measurements

The previous subsection showed how to evaluate the interference caused by transmissions in overlapping channels. Note that the spectrum densities studied correspond to activity periods, i.e. a snapshot taken when the transmitter is actually transmitting. It is therefore logical to draw the conclusion that a node that injects only a few frames per second is a source of interference that we can disregard in the presence of another node transmitting at the highest possible rate, even though the first station's frames are received with much more energy than the latter's. Therefore, the next step is to include the effect of utilization as a new parameter in order to consistently quantify interference from adjacent channels. Here, utilization is understood as the portion of time the node is actually transmitting into the air.

To do so, the mean power received from a source in saturation state is taken as a reference (considered 100% utilization[2]). Each utilization degree is measured with a spectrum analyzer, which covers 44 MHz around the center frequency f_c 10 times per second, taking samples every 100 kHz. Figure 3.13 shows the resulting average attenuation equivalent to a given degree of utilization. This equivalent attenuation is computed as the difference between the mean power received from a saturated station (i.e. no attenuation) and the mean power measured with lower utilizations.

[2] Note that the utilization taken as reference does not actually correspond to a busy time of 100% due to backoff intervals and inter-frame spaces

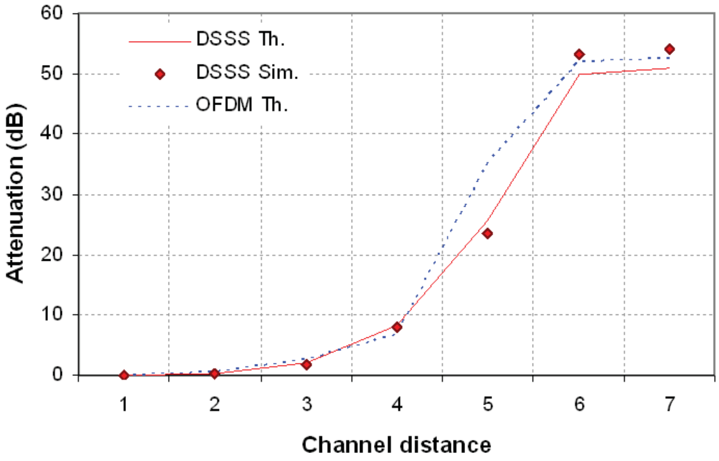

Figure 3.12: Attenuation for adjacent channels in IEEE 802.11 OFDM and DSSS PHYs

The values provided by the spectrum analyzer are directly proportional to the decay of the utilization. Thus, the attenuation in dB corresponds to the equation $10log_{10}(u)$, where u represents the utilization ($0 < u \leq 1$). That is, an utilization of 50% means that the source is transmitting half of the time, and hence the averaged measured power is halved (i.e. the equivalent attenuation is 3 dB).

Let us conclude our characterization of interference in partially overlapping channels with a complete example: if a receiver is tuned to channel 1, an interfering source continuously transmitting frames in channel 1 ($c = 0$) will not be filtered. But if the transmitter moves to channel 4 ($c = 3$) and reduces its utilization to 50%, the interference it produces is reduced on average by 8 (filter) + 3 (utilization) = 11 dB.

Although the utilization model's accuracy was validated with practical measurements with a spectrum analyzer, it is not applicable when the PER is computed from SINR values. The effects of utilization will depend not only on the ratio of active-to-idle periods, but also on the frame size distribution and the process describing inter-frame times in the interfering source. For example, whenever the filtered signal is above P_{th}, the physical carrier sense mechanism will defer transmission for all the interfering station's frames, so in this case the interference does not affect the PER and therefore there is no sense in applying the attenuation due to the utilization factor. The resulting effects of an adjacent channel interference above P_{th} are those of the co-channel interference. Otherwise, the interference can be added to the SINR before deriving PER. This observation is supported by the practical measurements carried out in [52].

The resulting PER can be approximated by $\text{PER}_u = u{\cdot}\text{PER}_1$, where PER_u is the PER obtained with an adjacent interferer's utilization of u and PER_1, the PER when the interferer's utilization is in saturation ($u = 1$). This PER variation can be translated to an equivalent $log_{10}(1/u)$ increase in the SINR.

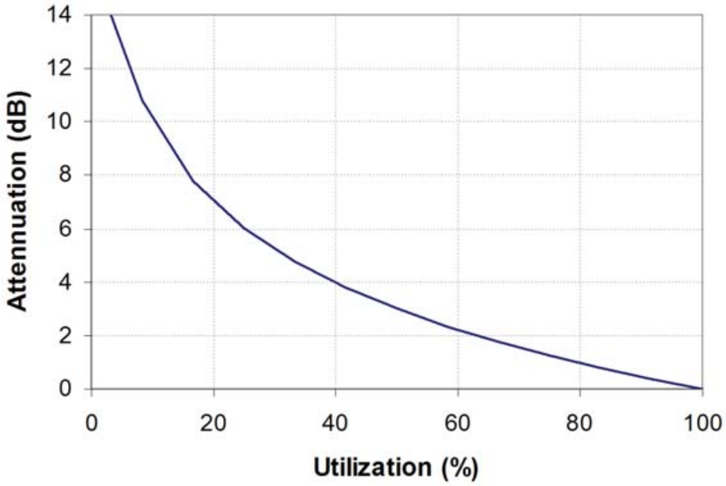

Figure 3.13: Average attenuation equivalent to a given utilization

Returning to the example used above, if a receiver is tuned to channel 1, an interfering source continuously transmitting frames in channel 1 ($c = 0$) will not be filtered. However, if the transmitter moves to channel 4 ($c = 3$) and reduces its utilization to 50%, the interference it produces, with regard to PER computation, is on average attenuated 8 (filter) + 0.3 (utilization) = 8.3 dB.

3.3.5 Co-channel interference

In order to promote coexistence among different technologies, devices using ISM bands must meet a number of international regulations that limit transmission power and force nodes to spread their signals. IEEE 802.11 goes beyond these regulations and adds other mechanisms to improve coexistence. Unfortunately, those mechanisms intended to politely accommodate other transmitters, such as carrier sense, also increase susceptibility to interference from other technologies.

As detailed in Section 2.3, the way the nodes of a WLAN share the medium is similar to an Ethernet segment. A CSMA/CA is used as medium access control scheme. Nodes sense the air interface before transmitting a frame; if it is busy, they will wait until it is released. This makes the study of interferences in IEEE 802.11 WLANs quite different from what is done in other radio networks due to the particular influence of interferences produced by cells using the same channel (co-channel interference): in a cell suffering only from co-channel interference, even though there is no traffic on it, the nodes may defer their transmissions if, when sensing the medium, they detect other nodes using the channel from an interfering cell. What is more, any transmission that is received with energy above P_{th} (see Section 2.3.2) will cause an IEEE 802.11 STA to defer its transmissions. This behavior entails an undesired side effect known as the *exposed node*

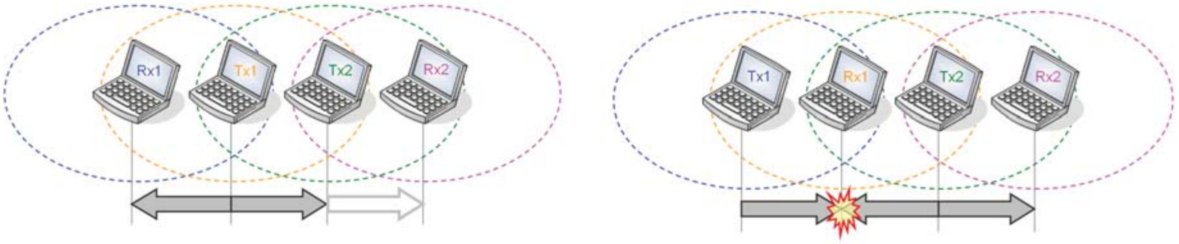

(a) Exposed node problem (b) Hidden node problem

Figure 3.14: Undesired effects in CSMA-based access networks

problem (c.f. [157]). This effect is exemplified in Fig. 3.14(a): T_x2 is willing to send a frame to R_x2 but it finds the channel busy due to an ongoing transmission between T_x1 and R_x1. Therefore, T_x2 will defer its transmission even though R_x2 is able to receive that frame, since it is not interfered by T_x1's transmissions.

Another known problem caused by co-channel interference in CSMA networks is the *hidden node problem* [276]. This undesired effect is represented in Fig. 3.14(b): T_x1 is transmitting to R_x1. During T_x1's transmission, a new frame arrives at T_x2's transmission queue, destined to either R_x1 or R_x2. This time T_x2 finds the channel idle (T_x2 cannot hear T_x1's transmissions). Therefore, both T_x1 and T_x2's frames collide at R_x1 and consequently both frames may get lost.

Both hidden and exposed node problems may be alleviated by means of the four-way handshake described in Section 2.3.3. To avoid the hidden node problem, T_x1 will send an RTS frame to R_x1. If the RTS frame is successfully received, this request is responded with a CTS. T_x2, which is in the transmission range of R_x1, "hears" the CTS and will therefore defer any transmission during the time specified in the CTS.
In the case of the exposed node problem, a STA hearing an RTS but not the corresponding CTS, can deduce that it is an exposed node and is permitted to transmit [68] to other nodes (e.g. T_x2 to Rx_2). Nevertheless, this requires an extra synchronization between the transmitters: if both T_x1 and T_x2 transmit at the same time, avoiding the exposed node problem, control frames sent by R_x1 and R_x2 (e.g. a CTS or an ACK) may not be received correctly.
Different approaches have been proposed to cope with the undesired effects of co-channel interference. Basically, they try to adapt the RTS/CTS mechanism dynamically in order to provide better MAC throughput efficiency (e.g. see [202] and [171]).

However, unlike other wireless technologies, co-channel interference in IEEE 802.11 WLANs may become less harmful than adjacent channel interference. A simple experiment can be carried out to illustrate this statement: the scenario consists of two totally overlapping WLAN cells, each composed of one 802.11b AP and one client station. One of the clients, used as the interference source, is pushed to saturation (i.e. there is a frame

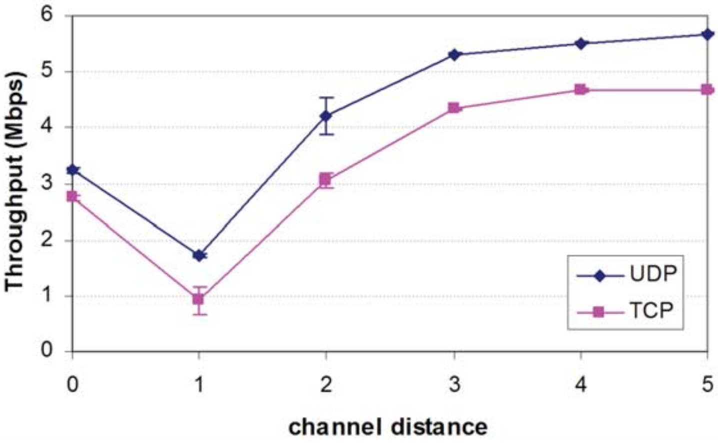

Figure 3.15: Throughput with interference from various channels (conf. interv. for 95%)

ready to be sent in the transmission buffer at all times). The other is used to test the UDP and TCP capacity of its cell for different channel settings: while one of the cells remains in channel 1, the other moves from 1 to 6. The throughput measurements are shown in Figure 3.15. Observe that under co-channel interference the performance is better than that obtained with 1 channel separation, which shows a minimum. With co-channel interference (channel distance $c = 0$), both transmitters are able to share the medium due to the CSMA/CA access scheme. However, with channel distance > 0, transmissions from the other station are RF-filtered and the received signal may drop below the carrier sense threshold (P_{th}), so the transmitters will not be able to avoid the possible collision. In this case, simultaneous frames sent on overlapping channels are treated as noise, and the desired packet is received with errors; hence the throughput degradation. As we increase the channel distance, the RF filter discriminates interfering signals to the point where the interference is completely removed ($c > 4$). In consequence, we can conclude that it is sometimes preferable to use partially overlapping channels and sometimes not. Section 4.3 will shed more light on these issues.

3.4 Simulation environment

The use of network simulators is a common practice. Simulators are not only used by the networks research community, but also by network administrators. Network simulators serve a variety of needs. The first group use simulators to evaluate the performance of new protocols, or to test changes and extensions to the existing ones. Simulators are also particularly useful in allowing network administrators to detect and correct potential problems before deploying a complex (and possibly expensive) scenario. Compared to the cost and time involved in setting up an entire testbed containing multiple networking devices (e.g. computers, sensors, routers, hubs, APs, etc.), network simulators are relatively fast and inexpensive.

This section is used to present the tools employed to perform the simulation-based evaluations described in different sections of this thesis. After reviewing different approaches available today, this section also gives details about the SIPTE11 implementation, including the key assumptions and simplifications that are made.

3.4.1 Open-source simulators

Without any doubt, Network Simulator 2 (ns-2) [219] is the most popular simulation tool among network researchers [188]. Its success is in part a consequence of the open-source model (source code license under GPLv2). Thanks to this developing paradigm, many independent researchers have contributed with extensions and new features. Therefore it has arguably the most extensive model set for research on Internet protocols and algorithms.

The ns-1 emerged around 1989 as a variant of the REAL network simulator, which in turn was a modified version of NEST 2.5 [180]. Its development has been supported by various grants over the years. Currently, the ns-2 development and its graphical user interface (GUI), known as Network Animator (NAM), are still supported by the Defense Advanced Research Projects Agency (DARPA), the National Science Foundation (NSF) and the University of Southern California/Information Science Institute (USC/ISI).

The core engine of the discrete-event ns-2 simulator is written in C++ programming language. However, the simulations are configured by means of OTcl scripts, an object extension to Tool Command Language (Tcl) developed at the Massachusetts Institute of Technology (MIT). This adds extra complexity to the process of adding new extensions or modifying existing objects. Nevertheless, most of the IEEE 802.11-related modules have been contributed by developers not directly associated with the project, and hence these extensions are not included in the main ns-2 distribution.

For the MAC layer there have been five different contributions[3]. Unfortunately, these contributions have been developed separately and only one of them can be chosen. The first contributions implemented the IEEE 802.11 DCF access, but soon PCF was also included. After the standardization of IEEE 802.11e, several implementations of HCCA and EDCA appeared. Probably the most comprehensive module is that of INRIA-Planète Group [21], which includes a solid IEEE 802.11a PHY layer implementation. This contribution is the basis of the IEEE 802.11 module in YANS and ns-3 simulators.

Due to the many shortcomings of the previous ns-2, several developers are currently exploring a move to a new simulator core. The first lesson learned from ns-2 is to avoid a dual-language simulator in order to decrease the overall complexity of the system. In other words, new models should be devised in a single language, while it should be possible to use the simulator from any other language (Python, Perl, Tcl, etc.). With this philosophy,

[3] ns-2 contributed code list: http://nsnam.isi.edu/nsnam/index.php/Contributed_Code

a major revision of the Network Simulator is born, the new ns-3 [32]. The single-language core of ns-3 is written in C++; in this way the integration of C/C++ modules adopted previously by ns-2 is straightforward. Nevertheless, ns-3 cannot be considered just a new version of ns-2, since it is being developed from scratch following a completely different architecture.
More precisely, ns-3 inherits part of its source from the solid-core event-based simulator YANS [115] code base. The basis of the YANS, detailed in [193], include emulation capacity, integration of user/kernel-space networking protocols or applications, and easy packet delivery tracing by means of widely-accepted formats (e.g. *pcap* for packet traces). The ns-3 project began development as of July 1, 2006 and is expected to take four years.

The concept of simulator/emulator and the integration with the OS kernel/user-space IP stack announced for the ns-3 is not new. As a matter of fact, this simulation methodology was baptized by S-Y. Wang as kernel re-entering [287]. This is the basis of the NCTUns simulator [290]: by using the mentioned kernel re-entering technique, a real-life Linux kernel's protocol stack is directly used to run the simulations. Besides, it can run any real-life UNIX user-space application program on a simulated node without any modification, and the commonly used Linux networking tools and commands (e.g. *traceroute*, *ifconfig*, etc.) can be run on a simulated network to configure or monitor the simulated network.
NCTUns implements IEEE 802.11b PHY and DCF, EDCA and HCCA MACs. It also allows multi-channel, multi-interface nodes for wireless mesh networks.

The OMNeT++ discrete event simulation system [33] has become quite popular recently. Components (modules) are programmed in C++, then assembled into larger components and models using a high-level language (NED). It is free for academic and non-profit use; commercial users must obtain a license.
In fact, OMNeT++ is not a network simulator per se; its generic nature served to a variety of different simulations, among them, of course, network simulations through the INET and Mobility Frameworks. As happened with ns-2, different projects developed different IEEE 802.11 modules uncoordinatedly. None of those modules provide a complete implementation of the IEEE 802.11 MAC and PHY layers.

GloMoSim [6] is the simulation tool developed at University of California at Los Angeles (UCLA) and is freely available for educational institutions. It is a scalable simulation environment for both wired and wireless large-scale networks since it allows parallel and distributed execution on a set of parallel computers. The parallel discrete-event simulation capability is provided by the use of PARSEC, a C-based simulation language developed by the same group.
The 802.11 modules are very limited and only offer ad-hoc DCF access. However, the main drawback is due to the fact that GloMoSim is no longer supported in the context of the original project. Since its last release in December 2000, GloMoSim has become a commercial project under the name of QualNet [34]. QualNet provides an extensive suite of implementations of models and protocols and expands GloMoSim Capabilities

for scenario visualization and analysis. With regard to IEEE 802.11, it provides the QualNet Wi-Fi, a specialized library of protocol models designed for WLANs, including infrastructure mode support, PCF, etc.

Another popular simulation environment is the Opnet Technologies Inc.'s Modeler [1]. Opnet Modeler is a commercial software, but it is freely available for academic use through Opnet's University Program. However, the academic version used to be limited in features. Nowadays they offer a full-featured Opnet Modeler for teaching purposes which does not include the Wireless Suite [2].

Opnet Modeler includes contributed packages from the academia. These and other modules are open-source and can therefore be modified and extended. The deployment is done through the GUI; the lowest level models must be structured as finite state machines and the specification of each state is coded using a C-like programming language called Proto-C.

Per contra, the simulator core source is not in the public domain. Moreover, comparisons performed between ns-2 and Opnet Modeler [118, 119, 207] conclude that both tools provide similar results for wired and wireless simulations[4]. Opnet, in contrast to the free open-source solutions, offers technical support (non-free). For these reasons, Opnet is preferred by the industry world while ns-2 is the most popular in the academic research field.

3.4.2 SIPTE11 simulator

SIPTE11 was developed at the Telematics Department of the UPC. The original code was created to simulate a pre-standard WLAN technology in cooperation with Alcatel. By the year 2000 it was revised and extended as part of the project CICYT TIC2000-1041-3-1[5]. The simulator was modified and updated by E. López Aguilera as part of her PhD Thesis [208] in order to fullfill the requirements of the IEEE 802.11 [3].

This discrete-time network simulator is written in C++ language and follows all IEEE 802.11 MAC protocol details. An evaluation and some implementation details of this simulator were presented in [209]. The tool allows the study of different scenarios at the link layer level. This includes different topologies: ad-hoc or infrastructure-based, single or multicell scenarios. It enables the evaluation of different PHY layers: IR, FHSS, DSSS or OFDM. Finally, all IEEE 802.11 MAC access mechanisms are also available: DCF (including both basic and RTS/CTS mechanism) and PCF.

Its modular design allows the substitution of the entire IEEE 802.11 MAC by a different access technology. Besides a module implementing the MAC functions, there is a module controlling all node events (management of reception/transmission queues), another module controls the events at a BSS level, and finally a single instance of the system module controls the simulation core process.

[4] Comparisons do not cover the Opnet Wierless Suite

[5] Funded by the Spanish Ministry of Education and Science

The following sub-sections are intended to provide interesting details about the operation of the simulator, paying special attention to the treatment of interference, and also describing the scenario generation process.

3.4.2.1 Interference Management

At this point we need to recall that this thesis is focused on the study of radio resource management (RRM) mechanisms. For this reason it is required that the tools employed to evaluate those mechanisms treat all issues related to interference in a proper way.

As seen in Section 3.3.4, the presence of adjacent channel interference produces an increase in the number of packets received in error. This effect is implemented in the simulated packet reception process as follows.
In reception, a frame is split up into intervals. For each of these intervals, the average SINR is computed as the quotient between the signal power of the desired frame and the sum of thermal noise power and interference signals. Any transmission occurred in a different cell during the desired frame reception will increase that interference (see example in Fig. 3.16) given that the interfering energy received is below the CS threshold. Otherwise, the transmission will be computed as a collision. The energy of a packet transmitted in a different cell will be affected by the filtering factor, which depends on the channel distance between the transmitter's and the receiver's frequency (see Table 3.2). After all these considerations, the BER is obtained for each separated interval. Recall that the BER is a function of the SINR (see Section 3.2). Then, the PER is derived from the BER and the interval length in bits, according to eq. 3.16. Finally and for each interval, a random value between 0 and 1 is drawn from a uniform distribution before deciding whether the packet is successfully received or not. If any of the random values happens to be smaller than the PER of its corresponding interval, the packet is assumed to be erroneous. Otherwise, the packet is successfully received.

Furthermore, co-channel interference requires a different treatment. As stated in Section 3.3.5, co-channel interference produces the potentially harmful effects known as the *hidden* and *exposed node* problems. Note that co-channel interference is not limited to the interference produced by nodes using the same frequency channel. We also consider as co-channel, all interference from partially overlapped channels, which is received with sufficient energy so as to produce the same effect as the strict co-channel interference. That is, co-channel is any transmission from a partially (or totally) overlapping channel that, even after filtering, is received with an energy above the carrier sense threshold (P_{th}). In consequence, this interference will effectively prevent the utilization of the medium due to the CSMA/CA mechanism. If one of these co-channel transmissions occurs during the reception of the desired frame, that frame will suffer a collision and it will not be successfully received. This property makes a "collision as failure" assumption, i.e. a packet collision results in a packet corruption. Exploiting of the *capture effect* is therefore not enabled [101] in our simulations. The capture effect, also called *co-channel interference*

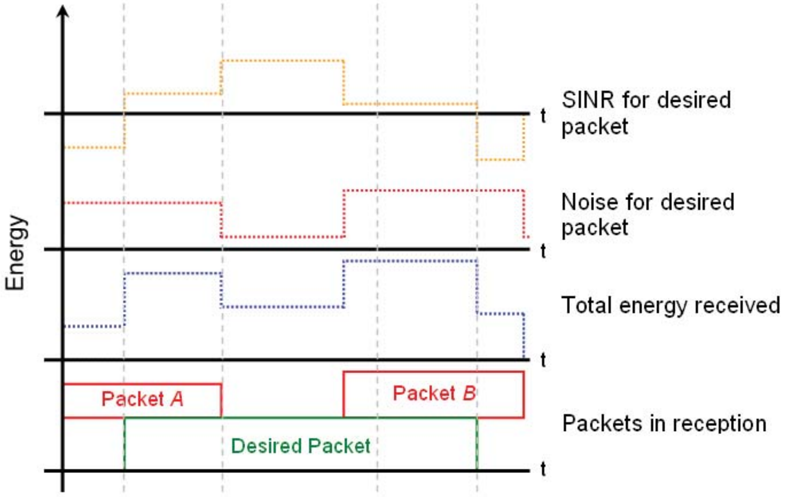

Figure 3.16: SINR vs time for a received packet

tolerance, is the ability of certain radios to receive a strong signal correctly from one transmitter despite significant interference from other transmitters [293].
As previously explained, and depicted in Fig. 3.16, the channel energy detected when a frame is sent will depend not only on its own cell utilization, but also on any ongoing transmission occurring in other cells.

3.4.2.2 Input/Output interfaces

The SIPTE11 simulation tool does not have a GUI. The configuration of the scenario is solely based on five different input text files. This characteristic could be regarded as a weakness, but is in fact an interesting feature. This type of input hampers the creation of individual and particular scenarios by the user. On the other hand, this structure enables the automatic creation and execution of uniform or random large-scale scenarios. This feature is especially interesting when the execution of a large number of simulations with varying scenarios is required. Of course, it requires the development of a new code intended to create the scenarios (i.e. intended to build the five different input files).
The required input files are listed and described next:

- *Ber.txt*: encloses the relation between SINR and BER for the different PHY modes enabled in the scenario. This way the simulator outsources the error model. The file is built according to the desired set of modulations (see Section 3.2).

- *Power.txt*: this file contains the connectivity or interference matrix. The signal power values at reception (before filtering) between each pair of nodes is expressed in dB. These values depend on the propagation model used.

- *Delay.txt*: similar to the *Power.txt*, this file contains the delay matrix. That is, this file contains propagation delay values computed for each pair of nodes in the system.

- *System.txt*: is used to configure general PHY and MAC parameters that are common to the whole system. By tuning inter-frame space times, preamble and header lengths, control frame sizes, etc., this file allows the correct set of parameters that define the PHY layer (e.g. DSSS or OFDM).

- *LanX.txt*: (where X identifies a particular BSS within the system) defines a single cell. It includes the number of transmitters and their traffic profiles (frame size, inter-frame time and physical bitrate).

The output of the simulation is contained in two different files. *Results.txt* and *NodeZ.txt.* The first one provides a comprehensive set of statistics computed for the whole system, detailed in a per BSS fashion. It provides the offered traffic, aggregate throughput, collision probability, PER, and a long etcetera. The other output file, *NodeZ.txt* (where Z identifies a particular node within a BSS), provides the above statistics, but in this case as seen from that particular station's point of view.

3.4.2.3 Scenario generation

As mentioned before, the SIPTE11 simulation tool outsources the error and propagation models. For these reasons, it is required to perform a pre-calculation of different radio parameters before the simulation is run.

In order to perform the simulation-based evaluation of the different scenarios and to facilitate the execution of thousands of simulations, a companion software (*EscGen*) is developed in this thesis, with the intention of creating new multicell scenarios automatically. The generated scenario is based on the uniform or random placement of a configured number of APs over a flat square terrain. Then, a given number of stations are randomly placed in the scenario. The x, y coordinates can be chosen either by following a discrete uniform probability distribution or by defining a smaller area where a given portion of the stations are concentrated so that hotspot-like scenarios can be emulated.

Once all nodes are placed in the scenario, the delay and the received signal power are computed for each pair of nodes, according to the chosen propagation model. This computation is used to create the *Delay.txt* and *Power.txt* input files.

There is a default association and rate selection scheme, but different criteria may be used. By default, stations are associated to APs based on received signal strength. The physical rate that each node uses to send data is chosen assuming a perfect adaptation, i.e. the modulation yielding the better performance is used by default (see 4.3.3). The user's offered traffic can be set at random or by selecting three predefined profiles (*low*, *medium* or *saturation*). This configuration is used to create *Lan.txt* and *System.txt* input files.

The main drawback of the SIPTE11 simulator is that it only considers a static scenario during the execution. That is, nodes are not mobile, traffic profiles remain unchanged during the simulation, and what is more, the SNR and physical rate are constant. In this way, realistic propagation issues (including rapid and slow fading) cannot be treated properly. In order to overcome this serious flaw, the EscGen is able to add those changes in the scenario between consecutive SIPTE11 executions. In other words, EscGen defines the initial scenario and calls SIPTE11; the simulation output is analyzed and the scenario is modified according to a realistic propagation model, varying traffic demands, moving stations, enabled RRM mechanisms, etc. Next, SIPTE11 is called again to simulate the new scenario, again the new output is analyzed, the scenario is modified and so forth.

3.4.3 Common assumptions

After a comprehensive survey of wireless research papers, Kotz et. al. came to the conclusion that the many assumptions and “axioms” made by researchers with regard to their simulations were far from reality. In [187, 188] those “axioms” are enumerated and later refuted by means of practical measurements. The list of axioms is as follows:

0. The world is flat
1. A radio’s transmission area is circular
2. All radios have equal range
3. If I can hear you, you can hear me (symmetry)
4. If I can hear you at all, I can hear you perfectly
5. Signal strength is a simple function of distance
6. Each packet is transmitted at the same bitrate

Obviously, the real world is not flat. Even at short distances, hills and buildings present obstacles that dramatically affect wireless signal propagation. An accurate 3D terrain model is essential for predicting the network performance in a particular and concrete scenario. It is important to point out that in our simulations we have considered a 2D scenario, since we were not focused on any specific location. However, the presence of obstacles at arbitrary locations within the scenarios may be included by adding the effects of shadowing in the propagation model. The optional Opnet’s Wireless Suite along with the Terrain Modeling Module, as well as QualNet, offer the possibility to introduce depth contour terrain maps in order to predict the effects of scattering, diffraction, etc. in the wanted region.

Measurements state that all radios are not identical: even though the same vendor and model is used for all the nodes’ radio of an experiment, in practice the radios and/or

antennas vary from node to node. This also leads to the conclusion that antennas are not perfectly omnidirectional, and thus, adding the effect of obstacles, it entails that the radio's transmission areas are not circular. Therefore, axioms 1 and 2 are not met. We have to admit that we adhere to these two axioms, yet only in part. Again, we do not need to simulate any particular network, with a given specific hardware. For these reasons we can assume "generic" devices with identical behavior. However, as mentioned before, a perfect circular coverage is withdrawn by adding random variations in the propagation model. This attribute of the propagation model also avoids axioms 3 (symmetry) and 5 (received power only depends on distance). Nevertheless, the shadowing model we adopted does not consider correlations: a real shadowing effect has strong correlations between two locations that are close to each other. Note that the shadowing effect should also be temporally correlated
The ns-2 simulator includes a free space model, a two-ray ground model, and also adds a log-normal non-correlated "shadowing" model. Similar to eq. 3.7, this model adds a zero-mean Gaussian distributed random value to the received power so as to account for the fact that surrounding environmental clutter can be very different at various locations at the same distance. As part of the work presented in [207], we found that using the ns-2 shadowing model with exponent 2.7 and 4 dB of standard deviation, the relative error with respect to the measurements presented in Section 3.1.5 were below 8%. Rayleigh/Rician fading is also available as a separate module [249] (also available for Opnet Modeler). YANS and ns-3 inherite ns-2's propagation models and natively add slow flat fading channels supporting both Rayleigh and Rician cases. NCTUns also reutilizes the ns-2 PHY layer code but adds a Rayleigh fading model
OMNET++ only provides the free-space propagation model, while GloMoSim adds the Two-Ray large-scale path loss models and an implementation of Rician Fading channel.

Many simulators assume axiom 4 because of their simplified error model. This model assumes that the packet reception probability distribution over distance (or over SINR) exhibits a sharp cliff, that is, the PER is either 0 or 1. For example, both ns-2 and Opnet common models assume that a transmission is perfect within the transmission range of a radio (as long as there are no collisions). When the borders of the coverage area are crossed, the PER suddenly raises to 1. Although both simulators allow the use of certain hooks to add more precise error models, these hooks are not commonly used. This issue is revised in the NCTUns simulator, which includes BER vs SINR curves as an easy-to-use optional channel. Qualnet and the GloMoSim project implement more sophisticated models. The approach followed by the YANS and ns-3 PHY implementations [181] are similar to the model used in our simulations: the PER is described as a function of the SINR (see Section 3.2). We have to stress that none of the mentioned simulators perform a real simulation of the PHY layer, but rather apply analytical models. Opnet Modeler in conjunction with the Wireless Suite does simulate at the lowest level[6], implementing the different modulations (DPSK, PSK, CCK and QAM) specified by the IEEE 802.11

[6] c.f. http://www.opnet.com/support/des_model_library/WLAN80211.html

standards.

As explained in the previous subsection, SIPTE11 also assumes axiom 6, that is, all frames are sent using the same PHY mode. The problem is that real devices perform some kind of rate adaptation in order to obtain the best possible performance according to the current signal quality (a more detailed discussion on this topic can be found in Section 4.3.3). However, our simulations were run using the EscGen scenario generator, which allows a different selection of the PHY mode as a function of the SINR. Excluding the new ns-3, the rest of the simulators assume a fixed rate, although some patches can be found for ns-2 that address these issues [117]. The ns-3 includes the AARF [194] rate selection algorithm by default.

Although not mentioned by Kotz in [187,188], I would suggest adding another axiom: *Two different cells utilize two independent channels*. In other words, many simulations consider that two cells using different channels do not suffer interference from each other. On the other hand, in ns-2 there is a single channel object that is common to all nodes and hence, all nodes belong to the same CSMA domain. The enhanced Network Simulator (TeNS) is an extension of the ns-2 that addresses these deficiencies (and others) [191]. According to [58], there is another ns-2 extension called ns-MIRACLE (Multi-InteRfAce Cross-Layer Extension) that implements adjacent-channel interference following the model we developed in Section 3.3.4. Opnet models different channels for different BSS, but the interaction between them in the default package is not properly implemented, since it only considers co-channel interference: if channels are partially overlapped, there will be collisions regardless of the channel distance; if channels are not overlapping, there will be no interference at all. However, packets received under the sensitivity threshold are added to the noise [239]. NCTUns alleguedly provides reliable wireles simulations [288], but the interaction between cells is not clear. YANS and ns-3 consider interference by following an approach that is similar to that explained for SIPTE11 in Section 3.4.2.1. For IEEE 802.11 PHY and MAC implementation details in these and other simulators, see [182].

To sum up, this section is intended to highlight the restrictions of simulation-based studies, which inevitably incur in unrealistic assumptions. However, this does not mean that simulations are not useful to illustrate the advantages or improvements of a given RRM algorithm, but the reader should be aware that those improvements shown by the simulations may be noticeably reduced in the real world. It is also worth to mention that this awareness pushed us to find ways of avoiding most of the axioms, as explained above.

3.4.4 Conclusions

In the preparation of this thesis we had the oportunuity to test different simulation tools by ourselves; among them, ns-2, Opnet, NCTUns and lately, ns-3. When it came to choosing the simulation environment for our evaluations, none of those tools fullfilled our requirements. The main drawback of all those simulators was the excessively simplified

PHY layers. For this reason, we chose the SIPTE11 tool. By outsourcing the error and propagation models through different configuration files, these models can be customized without the need of changing a single line of code. Furthermore, its interface based on input configuration files facilitates the execution of automated large-scale simulations. The presence of a GUI (e.g. Opnet or NCTUns) enables a user-friendly design of particular scenarios, but does not help when thousands of different scenarios are required.

Another criterion is related to licensing terms and availability of the source code. This discarded commercial solutions, such as QualNeT, but still left the door open for Opnet. The problem was that Opnet only offered a feature-limited version of the Modeler and did not include the Wireless Suite, in contrast to what they offer today. These facts led us again to SIPTE11. Given that this software was developed at this author's research group, we enjoyed an excellent "customer support" from the original developers, who very kindly helped us when a bug was detected or a new feature required.

However, if the choice of a simulation environment had to be made now, it would probably be different. The release of a stable and complete version of ns-3 will establish in the near future (at the moment of writing) a new standard among network researchers. One can easily predict that ns-3 will inherit the success of its predecessor ns-2, and will represent the new long-lasting architecture that will survive many years of abuse by both experienced researchers and newbie students. Opnet Modeler has also improved in recent years. Besides, nowadays both a complete version of the Modeler and the aforementioned Wireless Suite are available free of charge for academic research as part of the University Program.

Chapter 4

Throughput and capacity models

In order to evaluate the quality of a design prior to its deployment, we need tools capable of predicting WLAN performance. Simulators are often used, but they are expensive and the way they model interference in multicell environments is not accurate due to simplifications in the physical layer: co-channel and adjacent channel interference emulation does not behave as observed in real testbeds [188]. One alternative is to use analytical models, which also make significant assumptions. However, they are faster in providing a reliable approximation of a network's performance and the results are precise enough to allow the objective evaluation of various WLAN designs. Capacity estimations can then be used as the metric to compare different frequency channel schemes, different power allocations, user-AP associations, etc. In other words, capacity estimations can assist many RRM mechanisms that are in charge of dynamically tuning different WLAN parameters in order to optimize its performance.

The capacity limits of IEEE 802.11 access networks have been widely studied in the literature. These limits were first explored for a single user transmission, but subsequently efforts were concentrated towards modeling the capacity of a multi-user cell. Nevertheless, capacity of a multicell WLAN, which involves new phenomena such as inter-cell interference (co-channel and adjacent channel), has rarely been studied. Moreover, the results of these studies were either too simple (e.g. applying an independent single-cell model for each cell) or too complex to allow frequent and rapid estimations. Our aim is therefore to develop a straightforward algorithm that is able to provide capacity estimations in a timely manner so that they can be used by different RRM mechanisms as an immediate measure of the potential performance of a WLAN. The resulting model, published in [130, 135], takes both co-channel and adjacent channel interference into account, and also comprises the effects of rate adaptation.

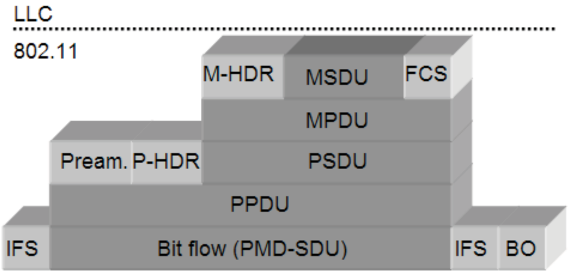

Figure 4.1: Overhead in IEEE 802.11 sublayers

4.1 Single user

The theoretical maximum throughput that is available to an 802.11 STA can be easily derived from the parameters specified by the different standards. As shown in [172], if we know the length of the different headers added by the different sublayers (see Fig. 4.1), we can compute the theoretical maximum throughput (TMT) as the maximum number of MSDUs that can be transmitted per time unit.
For this upper bound computation, the following assumptions are made:

- There are neither errors (BER =0) nor collisions.
- The transmitter always has data to send (saturation).
- Fragmentation is not used.
- Management frames (e.g. *Beacon*) are not considered.

At this point, it is required that we recall the frame sequence given for the transmission of an MSDU. In Section 2.3.3 two access methods were described: basic access and four-way handshake (RTS/CTS). These two methods involve a slightly different message exchange (see Fig. 4.2) and are briefly described here.

Using the basic CSMA/CA access, a STA first senses the medium and transmits if it is idle during a DIFS time plus a random backoff time (BO); the latter is used to avoid collisions with other STAs waiting to transmit. After the data frame is sent, the receiver waits for a SIFS time before sending and acknowledgment.
The RTS/CTS access was intended to reduce collisions where the hidden node problem [98] is likely to appear. Previous to the data transmission, an RTS frame is sent to the receiver, which will respond with a CTS announcing to its neighborhood that the medium will be occupied, thus avoiding collisions with nodes that are hidden to the transmitter. This mechanism improves the network performance in the presence of hidden nodes, but adds considerable overhead.

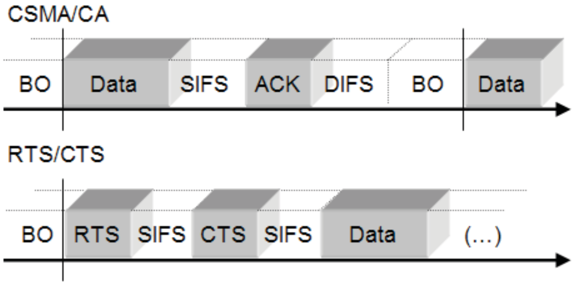

Figure 4.2: Message sequence for basic access and RTS/CTS

Then, the TMT is expressed as follows:

$$TMT = \frac{MSDU_{size}}{Time\ per\ MSDU} \tag{4.1}$$

$$TMT = \frac{MSDU_{size}}{T_{DIFS} + T_{BO} + T_{RTS} + 3T_{SIFS} + T_{CTS} + T_{Data} + T_{ACK}} \tag{4.2}$$

where T_{Data} is the time during which the actual data frame is transmitted (including preamble and headers). Note that this time changes depending on the modulation or the PHY rate used. IFS times and the duration of RTS/CTS and ACK frames are independent of the modulation since they are always sent at the lowest rate. Also note that eq. 4.2 represents the TMT for an RTS/CTS access. For the basic access, we have to subtract T_{RTS}, T_{CTS} and $2T_{SIFS}$ from the denominator.
Figure 4.3 shows TMT values for IEEE 802.11a/g and 11b modulations. A typical WLAN connected to an Ethernet segment renders a maximum throughput around 6 Mbps (11 Mbps CCK and basic access and 1500-Byte MSDUs).

This simple model can be improved by capturing the effects of transmission errors. This way we add a dependency on SINR. In [104] the ideas of [172] are adapted to include the extra delay caused by retransmissions that was not present in the TMT model.
Knowing BER values (see Section 3.2) for a given modulation M, P_b^M, the PER of an l-bits packet can be found according to eq. 3.16. Thus, the probability that a transmission is successful is as follows:

$$P_u^m(l) = (1 - P_{eDATA}^m(l))(1 - P_{eACK}^m) \approx (1 - P_{eDATA}^m(l)) \tag{4.3}$$

That is to say, a successful transmission requires that neither the data packet nor the ACK are received with errors. Given that the ACK frame is shorter than data frames and that it is usually sent using the most reliable modulation, the second term in eq. 4.3 can be neglected. Then, the probability of a successful transmission at the i-th attempt is obtained according to:

$$P_{succ}[i] = (1 - P_{succ}^m(l))^{i-1} P_{succ}^m(l) \tag{4.4}$$

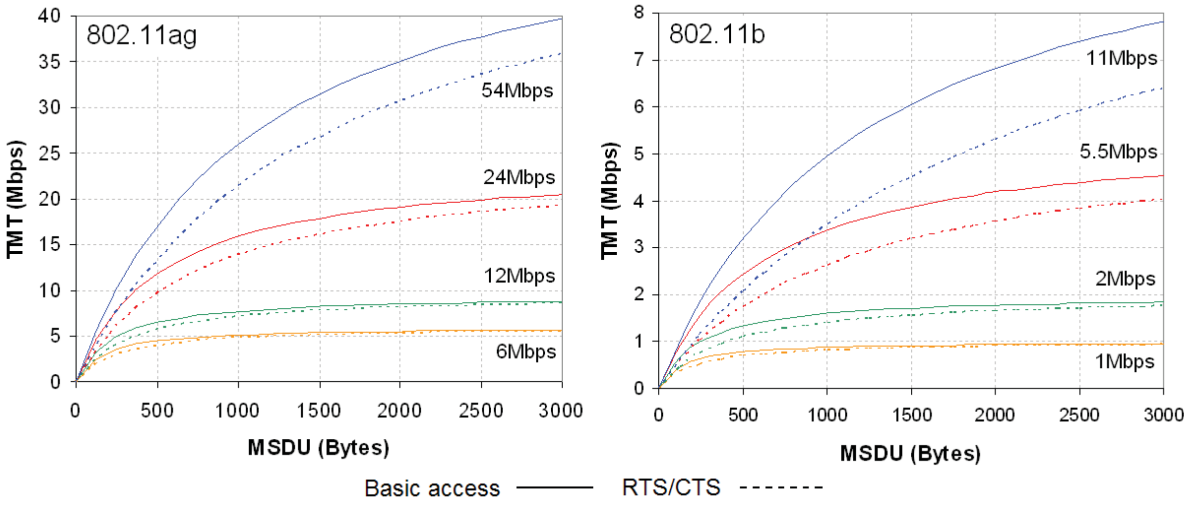

Figure 4.3: TMT vs. MSDU size for basic and RTS/CTS access and different 11a/g and 11b modulations

For each retransmission, the range of values that can be given to T_{BO} is doubled until CW_{max} is reached. Recall that CW_{max} defines the maximum number of time slots that a STA will wait in the backoff stage. The average backoff time after i consecutive transmissions of the same frame, $T_{BO}(i)$ is:

$$T_{BO}(i) = \begin{cases} \frac{2^i(CW_{min}+1)-1}{2} T_{slot} & 0 \leq i < m \\ \frac{CW_{max}}{2} T_{slot} & i \geq m \end{cases} \tag{4.5}$$

where m is known as the "maximum backoff stage" and together with CW_{min}, CW_{max} and T_{slot}, are all standard-dependent parameters. These values are summarized in Table 4.1.

Taking into account eqs. 4.4 and 4.5, and an infinite number of possible retransmissions, the *time per MSDU* that includes the effects of packet errors and retransmissions used in eq.4.1 is now:

$$Time\ per\ MSDU = \sum_{j=2}^{\infty} \left(P_{succ}[j] \sum_{i=1}^{j-1} T_{frame}(l, M, i) \right) + T_{frame}(l, M, 0) \tag{4.6}$$

where $T_{frame}(l, M, i)$ is the time needed to successfully send an $l-$bit frame using the modulation M. With basic access, $T_{frame}(l, M, i) = T_{DIFS} + T_{BO}(i) + T_{Data}(l, M) + T_{SIFS} + T_{ACK}(M)$.

In [252], the assumption of infinite retries is removed, given that in practice, after 4 or 7 retransmissions (depending on the magnitude of l), the packet is discarded.

	DSSS - CCK				**OFDM**			
Rate (Mbps)	**1**	**2**	**5.5**	**11**	**6**	**12**	**24**	**54**
T_{ACK} (μs)	304	248	212	202	50	38	34	30
T_{RTS} (μs)	352	272	221	207	58	42	34	30
T_{CTS} (μs)	304	248	212	202	50	38	34	30
$T_{preamble}$ (μs)	192		192/96		20			
T_{SIFS} (μs)	10				16			
T_{DIFS} (μs)	50				34			
T_{slot} (μs)	20				9			
CW_{min}	31				15			
CW_{max}	1023				1023			
m	6				7			

Table 4.1: IEEE 802.11 PHY parameters

Needless to say, that the literature on this topic offers a wider variety of approaches. For example, in [231] capacity predictions are based on previous empirical measurements. First of all, a data base is built with SNR/throughput relationships obtained in practical scenarios. Then, given an SNR value, a prediction of the throughput is provided according to the channel conditions. However, it does not take into account the effect of collisions or utilization. This idea is improved in [178] to provide per-link capacity estimations.

Note that these models are only valid to estimate the capacity of a communication between one transmitter and one receiver. In other words, the transmitter is not competing with other STAs in order to gain access to the medium. The capacity estimation in a BSS with multiple active STAs needs to capture the effect of collisions. Nevertheless, in [221] the TMT approach is used to compute the throughput of a cell with n STAs ($n > 1$). However, the authors themselves recognize that the complexity of the backoff mechanism in a BSS is not properly captured.

4.2 Single cell

Other models do include the study of collisions. In [79] and in [83] a *p-persistent* protocol for WLANs is analyzed. This protocol is similar (but not equal) to the IEEE 802.11 standard. In this case, the backoff interval is obtained by following a geometric distribution with probability parameter p.

Even though it makes several assumptions, the most popular approach used to model the saturation throughput of an IEEE 802.11 cell is authored by Bianchi [70]. This model

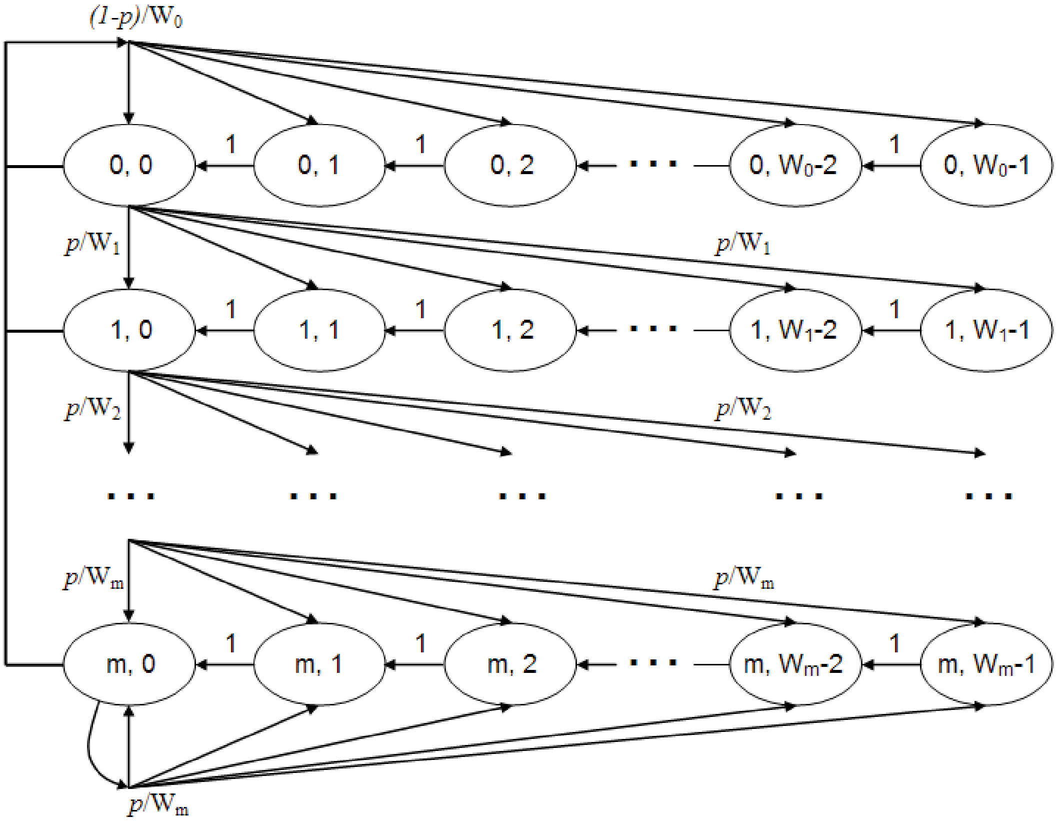

Figure 4.4: Bianchi's Markov chain model for the backoff window size

is based on the saturation condition, that is to say, all n nodes in a cell always have frames ready to be sent on their transmission queues. The n STAs are using the same PHY mode and the channel is considered ideal (i.e. no hidden nodes and no errors due to noise or interference). Furthermore, there is no coexistence between basic and RTS/CTS access.

The key concept of this model is that each packet injected into the medium collides with constant and independent probability p, regardless of the number of previous attempts. This probability p is considered a constant in the network under study. Let $b(t)$ be the stochastic process representing a STA's backoff timer. When $b(t)$ reaches 0, the STA is allowed to transmit. The constant m represents the "maximum backoff stage" so that $CW_{max} = 2^m CW_{min}$. For convenience $W = CW_{min}$ and $W_i = 2^i W$ where $i \in [0, m]$. Then, $s(t)$ is used to represent the stochastic process of the backoff stage (0, ..., m) for a given STA at time t. The backoff stage is increased from i to $i+1$ after an unsuccessful transmission.
Taking into account the previous assumptions, the bidimensional process $\{s(t), b(t)\}$ can be modeled with the discrete time Markov chain depicted in Fig. 4.4. In this Markov chain, the one-step transition probabilities are:

$$\begin{cases} P\{i,k|i,k+1\} = 1 & k \in [0, W_i - 2] \quad i \in [0,m] \\ P\{0,k|i,0\} = (1-p)/W_0 & k \in [0, W_0 - 1] \quad i \in [0,m] \\ P\{i,k|i-1,0\} = p/W_i & k \in [0, W_i - 1] \quad i \in [1,m] \\ P\{m,k|m,0\} = p/W_m & k \in [0, W_m - 1] \end{cases} \tag{4.7}$$

The first equation in 4.7 accounts for the fact that the backoff is decremented at the beginning of each slot time. It is important to note that a discrete and integer time scale is adopted: t and t+1 correspond to the beginning of two consecutive slot times and a slot time begins immediately following an idle DIFS. The time interval between two consecutive slot time beginnings vary, since it may or may not include a packet transmission. In other words, the interval between two consecutive slot times is the time between two backoff decrements.

The backoff stage is reset after a successful transmission, and hence the second equation. On the other hand, after an unsuccessful transmission, the backoff stage is increased (i.e. CW is doubled). This is represented by the third equation. Finally, the fourth case models the fact that once the backoff stage reaches m (i.e. $CW = CW_{max}$), it is not increased in subsequent packet transmissions.

Solving the Markov chain, the stationary distribution is represented by the probabilities $b_{i,k}$, with $b_{i,k} = lim_{t\to\infty} P\{s(t) = i, b(t) = k\}, i \in [0,m], k \in [0, W_i - 1]$. Then, the probability τ that a STA transmits a packet in a randomly chosen slot time is:

$$\tau(p) = \sum_{i=0}^{m} b_{i,0} = \frac{b_{0,0}}{1-p} = \frac{2(1-2p)}{(1-2p)(W+1) + pW(1-(2p)^m)} \tag{4.8}$$

The value of p is still unknown, but note that the probability p that a packet collides is the probability that, in a time slot, at least one of the n-1 remaining STAs transmits. Each remaining STA transmits with probability τ. Therefore:

$$p = 1 - (1-\tau)^{n-1} \longrightarrow \tau(p) = 1 - (1-p)^{\frac{1}{n-1}} \tag{4.9}$$

Equations 4.8 and 4.9 form a nonlinear equation system with a unique solution and with two unknowns, τ and p, which can be solved using numerical techniques. Once these values are known, other important relationships can be derived:

$$P_{tr} = 1 - (1-\tau)^n \tag{4.10}$$

$$P_{idle} = 1 - P_{tr} \tag{4.11}$$

$$P_{ok} = n\tau(1-\tau)^{n-1} \tag{4.12}$$

$$P_{col} = P_{tr} - P_{ok} \tag{4.13}$$

P_{tr} is the probability that there is at least one transmission in a slot time, and logically 1-P_{tr} (P_{idle}) is the probability that the slot time is empty. P_{ok} is the probability that there

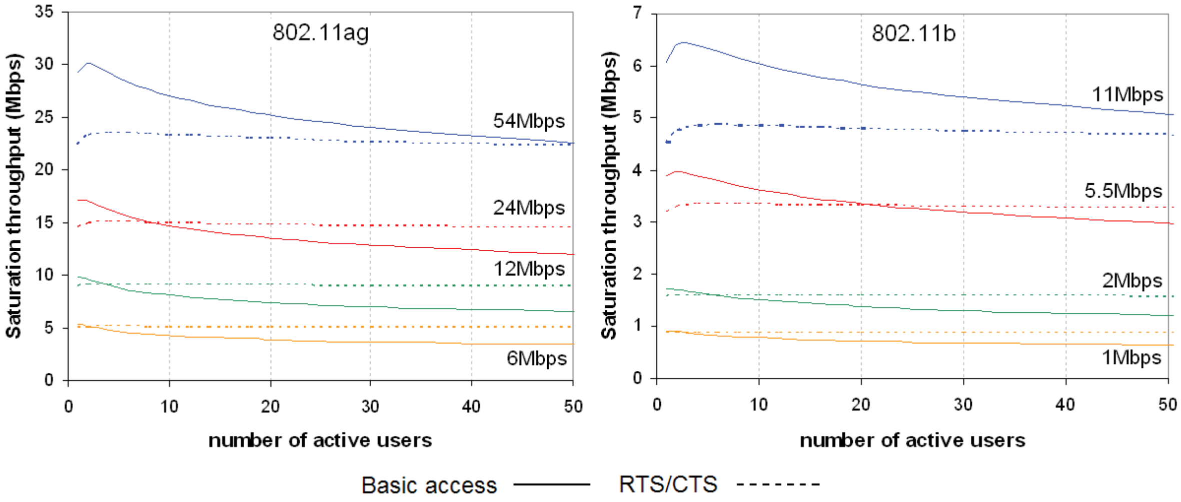

Figure 4.5: Saturation throughput in a cell vs. number of active users (MSDU size 1500B)

is a transmission and that this transmission is successful, and P_{col} is the probability that a transmission collides (at least one transmission, but none of them is successful).
Finally, the saturation throughput S can be defined as the average amount of payload bits that is successfully transmitted per slot time:

$$S = \frac{P_{ok}E[P]}{P_{idle}T_{slot} + P_{ok}T_{succ} + P_{col}T_{col}} \tag{4.14}$$

where $E[P]$ is the average payload size, T_{slot} is the duration of an empty slot, T_{succ}, is the previous *Time per MSDU* of eq. 4.2 subtracting T_{BO}, or more precisely, T_{succ} is the time needed to successfully deliver a frame. T_{col} is the time spent due to a collision: $T_{col} = T_{DIFS} + T_{frame}$ for basic access and $T_{col} = T_{DIFS} + T_{RTS}$ for an RTS/CTS access. Note that the time wasted because of a collision is smaller in the case of RTS/CTS access, since $T_{RTS} << T_{data}$. This difference provides a significant improvement when the collision probability is high (i.e. n is high), as shown in Fig. 4.5
Finally, note that the expression of S in eq. 4.14 is the same regardless of the PHY mode or the access method used. It is sufficient to substitute the constants (e.g. T_{DIFS}, T_{ACK}, etc.) for the values specified in the corresponding IEEE 802.11 standard (see Table 4.1).

Several papers have built on this basic model. The work in [294] and [189] extend Bianchi's analysis to include finite retransmission attempts providing a new expression for τ; [236] and [286] adapt the model to backoff mechanism variants. In [285], [90] and [274] the effect of error probability is introduced into the model. The constant probability p of eq. 4.8 and 4.9 is redefined. Now p is the probability that a packet encounters a collision or is received in error: $p = p_c + p_e - p_c p_e$. In the latter expression, p_c is the "old" collision probability and p_e is the error probability or PER. Then, eq. 4.9 is rewritten as:

$$p = 1 - (1 - p_e)(1 - \tau)^{n-1} \longrightarrow \tau(p) = 1 - \left(\frac{1 - p}{1 - p_e}\right)^{\frac{1}{n-1}} \tag{4.15}$$

Consequently, the average slot time is also modified in order to include the effect of packet losses, so that the new expression for the saturation throughput S in the presence of packet errors is as follows:

$$S = \frac{(1 - p_e)P_{ok}E[P]}{P_{idle}T_{slot} + (1 - p_e)P_{ok}T_{succ} + P_{col}T_{col} + P_{ok}T_e p_e} \tag{4.16}$$

The new parameter T_e is the average time that the medium is sensed busy due to an error transmission. For a basic access $T_e = T_{data} + T_{DIFS}$ and $T_e = T_{RTS} + T_{CTS} + T_{data} + T_{DIFS} + 2T_{SIFS}$ for an RTS/CTS access.

Finally, [164,165] includes the hidden terminal problem. In addition to the throughput analysis, some of the aforementioned papers also provide a companion derivation of the average transmission delay performance (e.g. see [85, 91]).
One step further, [301] and [84] revisited again Bianchi's model in order to obtain the performance in non-saturation conditions. Similarly, but following a more straightforward approach, in [114] two states are added to Bianchi's Markov chain to model unsaturated STAs. Moreover, in [84] and [114] the effect of multi-rate STAs is also included.

4.3 Multiple cells

Despite all the research done with the aim of modeling the capacity of the IEEE 802.11, only a few papers have focused on the multicellular scenario. The work in [113] is one of those few. However, each cell is treated individually, thus ignoring the evident effects of the interaction between neighboring cells. In this section we describe a new approach: our model is intended to provide capacity estimations for an ESS. The idea behind this model is to provide a metric whose optimization by means of an intelligent frequency management produces an improvement in the network performance. This algorithm should capture the effects of interference and should also provide new estimations rapidly so as to enable a rapid adaptation of the ESS to a dynamic environment.

For our evaluation, we perform a mathematical analysis based on Bianchi's model [70]. As detailed in Section 4.2, this model is used to compute the saturation throughput of an IEEE 802.11 WLAN for a single isolated cell that is free of transmission errors and hidden terminals; the number of competing stations is known and it is assumed that they always have frames to send. Adapting this model to a system based on multiple WLAN cells is not trivial, but we can apply several simplifications without deviating from the expected values. Two types of interference are considered: adjacent channel interference (caused by transmissions on overlapping channels); and co-channel interference (caused by nodes using the same channel). Moreover, most of the IEEE 802.11 implementations found on the market today apply proprietary link adaptation schemes in order to select the best modulation (physical rate) according to the current channel conditions. These algorithms are, in general, not available to the public, thus hampering the inclusion of rate adaptation in realistic analytical throughput models. For these reasons, in the resulting formulations we also introduce , the effect of a simple link adaptation algorithm. It should be pointed out that the heuristic we propose is intended only for modeling purposes and not for implementation. Our contribution, published in [130], is therefore twofold: to the best of our knowledge, we present the first analytical study of throughput performance including both types of interferences, co-channel and adjacent channel, and the effect of bit rate adaptation.

4.3.1 Effect of packet errors in saturation throughput

We depart from Chatzimisios' expression for saturation throughput in the presence of transmission errors [90]. The effect of errors is added to the model, given an average PER for a given cell; the PER can be derived from the SINR.

As explained in Section 3.3.4, we know that adjacent channel interference reduces the SINR, and hence reception errors appear. The PER for each node's transmission will depend on the SINR sensed by the corresponding AP upon reception of the node's frames. Recall that the SINR, and consequently the PER, also depend on the utilization factor u of the adjacent cells. Based on the SINR values and the modulation used by the station for its transmissions, we can derive the BER for a given client, which can be obtained theoretically using the formulas given in Section 3.2. However, to achieve more realistic results, empirical curves can be used. Usually, these BER vs. E_b/N_0 curves are included in the transceiver datasheet. For our tests, we use the data provided with Intersil Prism HFA3861B [9]. Once the BER is known, the PER can be approximated using eq. 3.16 ($PER = 1-(1-BER)^l$). We omit the effect of losing an ACK frame, since its probability is negligible due to its small size and the fact that it is usually transmitted at the slowest (i.e. most reliable) bitrate. We have to note that, as stated in [210], this approximation will not remain applicable for large values of BER.
Figures 4.6 and 4.7 show the saturation throughput for a single station in the presence of transmission errors. The figures are obtained analytically following [90], given a packet size of 1024 Bytes.

The effects of utilization in adjacent cells anticipated in Section 3.3.4 are verified by means of practical measurements and simulations. The simulated scenario consists of two partially overlapping IEEE 802.11a cells with two stations each (tx and rx). Cell A's rx receives -76 dBm of unfiltered signal from cell B's tx. The channel distance is 2 (ch. 34 and 36). B's utilization is increased from 0 to 100%, whereas A's tx is always in saturation state. Simulations are made with two different modes (18 and 54 Mbps)

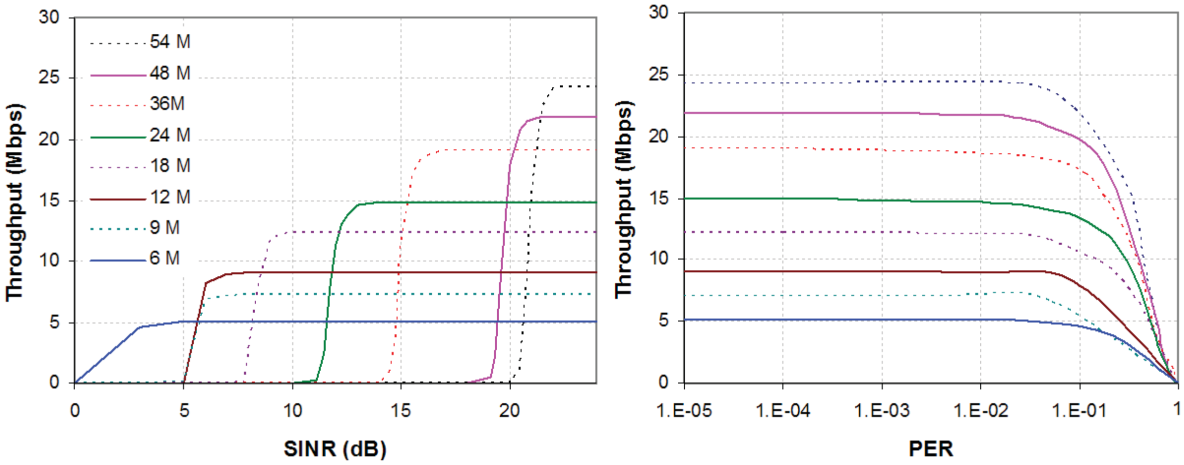

Figure 4.6: Effect of packet errors in an IEEE 802.11a/g transmission with 1024-Byte packets

in order to verify that our assumptions are independent of the modulation used. Details of the simulator's PHY layer implementation are given in Section 3.4.2. The results are shown in Fig. 4.8: dotted lines are obtained through analytical models with the saturation throughput computed using eq. 4.16 after deriving PER from u and SINR as previously explained; solid lines are obtained from simulations. Practical measurements are made with a similar testbed but using DSSS devices with 1Mbps PHY (DBPSK); PER and SINR show the same relationship.

4.3.2 Effect of co-channel interference

Due to the CSMA/CA mechanism, co-channel interference results in undesired effects, such as the hidden node and exposed node problems (cf. [157]). Bianchi's model was again revised to incorporate co-channel interference, but the authors of [241] only considered two interfering cells. In [140,141], a complete model is provided whose complexity hinders its implementation as part of a dynamic RRM mechanism. Our approach is based on the fact that WLAN throughput in the presence of co-channel interference can be modeled assuming long-term fairness among the nodes sharing the same channel. That is, in the long term, all stations have the same probability of gaining access to the common channel. Although this assumption ignores observable phenomena that arise in various topologies (e.g. see [93]), we find that the error introduced with individual flow calculations is not critical for the estimation of the global capacity, as shown in this section. In infrastructure mode, given a set of nodes N, the algorithm computes the available throughput for every

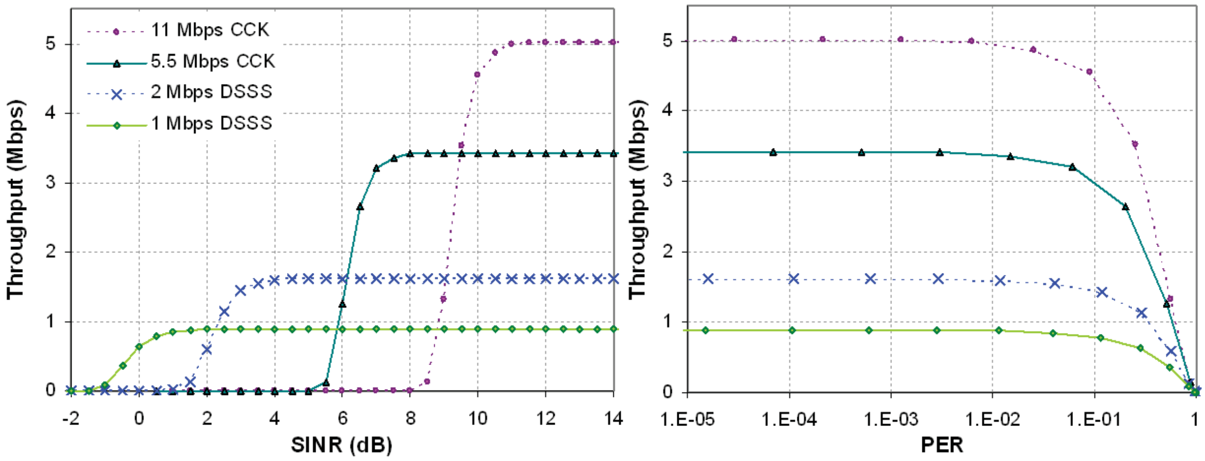

Figure 4.7: Effect of packet errors in an IEEE 802.11b transmission with 1024-Byte packets

node $i \in N$, S_i as:

$$S_i = s(i, C_i) S_{th}; \quad s(x, Y) = \begin{cases} \frac{1}{|Y|} & 1 - \sum_{j \in Y \neq x} u_j \leq \frac{1}{|Y|} \\ 1 - \sum_{j \in Y \neq x} u_j & otherwise \end{cases} \tag{4.17}$$

where $C_i \subset N$ is the subset of the elements of N that compete with STA i for the channel. This subset C_i depends not only on the nodes associated with i's AP, but also on the co-channel stations from other cells that are within carrier sense range of i. Thus, we include the influence of exposed nodes. The hidden node problem is minimized by using RTS/CTS. We find that the number of competing stations for i can be well approximated using the larger of the following two values: the number of nodes in i's cell and the maximum interference *clique* to which i belongs. Plotting stations and their interference as an undirected graph, a *clique* is a subset of nodes such that every pair of nodes are interferers. This method for modeling co-channel interference (henceforth *CL* model) involves the resolution of the *clique* problem, which is *NP-complete*. However, this is not an issue since a tree search function can be easily added to the code, which solves the *clique* problem without hindering the operation of the algorithm. If only the stations in range are taken into account (*CR* model), the resulting approximation is good for lower densities. Figure 4.9 shows the differences between *CL* and *CR*. The function $S_{th}(C_i)$ computes the saturation throughput for the set C_i, according to 4.16.

Finally, The multiplier $s(i, C_i)$ is used to reflect the effective share available to STA i. If the utilization of the remaining contenders is greater than $|C_i|^{-1}$, long-term fairness will ensure a share $s(j, C_i) = |Ci|^{-1} \forall j \in C_i$. Otherwise, the excess bandwidth not used by the members of C_i is available to i. Figure 4.9 illustrates this: following both models (*CL* and *CR*), node j will ideally receive 1/2 of the maximum throughput (ignoring collisions, errors and other effects to simplify the example), since there is only one competing station: i. However, i competes with 4 (*CL*) nodes, so at most it will receive just 1/4 of the share.

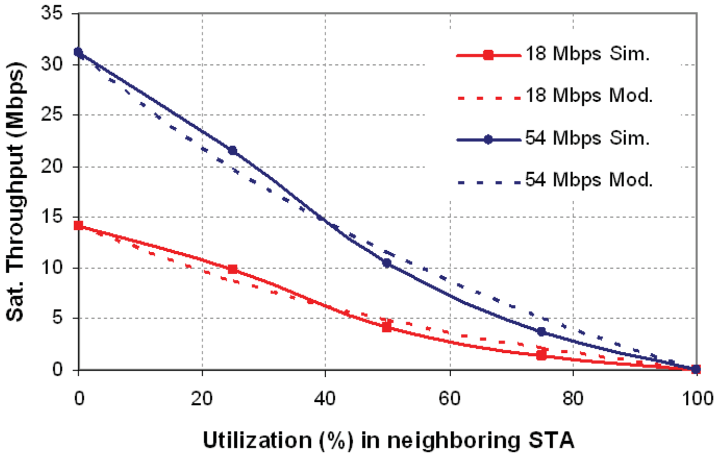

Figure 4.8: Saturation throughput with increasing utilization in adjacent cell

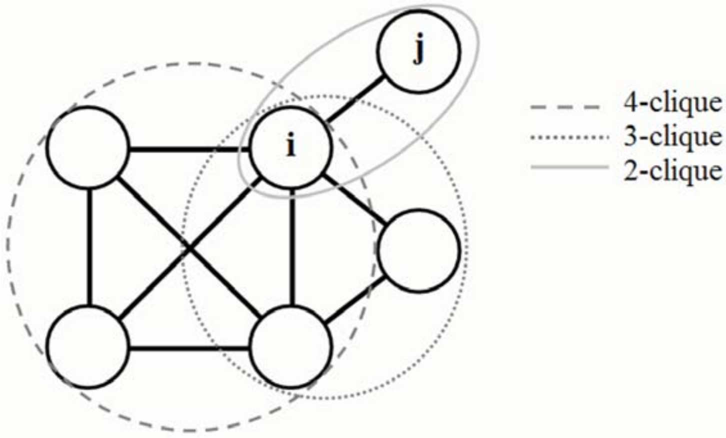

Figure 4.9: Interference graph and competing stations for node i 4 (CL) and 6 (CR)

Therefore, in practice, j is competing with just 1/4 of a node, thus obtaining 3/4 of the maximum throughput. Recall that u_j represents the node j's utilization (as defined in Section 3.3.4.1); in saturation, $u_j = S_j/S_{th}(C_j)$. In this case, the values $S_j \forall j \in C_i$ are required to obtain S_i. That is, in order to obtain S_i, we need to know the value of S_i in advance! To avoid this incongruity and minimize the error we introduce, we follow a simple heuristic: S_i is computed starting with the node with the greatest number of competitors (highest degree), and so on.

This co-channel model has been evaluated through extensive ns-2 simulations [219] (see Section 3.4.1). Note that wireless communication in ns-2 uses different and independent channel objects for different cells; as a result, cross-channel noise and interference are not simulated. To correct this, our simulation consists of ad hoc nodes using only one channel object. Four fixed stations playing the role of APs are located so that their coverage

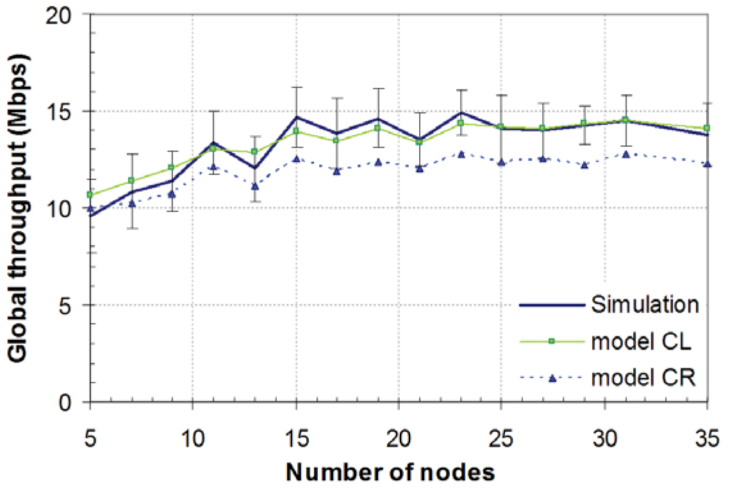

Figure 4.10: Evaluation through simulation of throughput models with co-channel interference

ranges partially overlap. The rest of the nodes are placed randomly throughout the area and configured to send data to the closest AP. All these 802.11b nodes are driven to saturation by setting up a constant bitrate (CBR) UDP source so that there is always a 1500 byte datagram ready for transmission in each node's queue. So as to isolate the effect of co-channel interference, no transmission errors are introduced. Figure 4.10 compares the results of the simulations for both models (*CL* and *CR*).

Our calculation of individual flows clearly introduces some error, given that we are using an average PER and an average PHY rate for the whole cell. As studied in [210], the case in which each station is subject to a different BER adds not inconsiderable complexity. Besides, the starvation problems measured in [92] are not taken into account. Notwithstanding the approximations applied on the computation of individual flows, Fig. 4.10 shows that the measure of the overall capacity is balanced, and that the error is reduced as the number of nodes is increased (especially for the *CL* model).
Also note that these approximations are only valid for infrastructure mode; ad hoc schemes usually require that nodes cooperate to forward each other's packets through the network. This means that the throughput available for each node's applications is limited not only by the raw channel capacity but also by the forwarding load imposed by distant nodes [166].

4.3.3 Modeling of a link adaptation algorithm

As seen in Section 2.4, IEEE 802.11 standards define several sets of modulations and coding rates for the different physical layers. For example, IEEE 802.11b specifies four modes: 11 Mbps (8-bit CCK), 5.5 Mbps (4-bit CCK), 2 Mbps (DQPSK) and 1 Mbps (DBPSK) to be used in the 2.4 GHz frequency band. Each different scheme provides a different transmission rate, but the higher the chosen rate, the worse it performs in the presence of noise and interference; i.e. given certain channel conditions, there is an optimal modulation that maximizes the throughput. It is interesting to observe that the modulation selection does not depend on the frame size: big frames render a higher throughput than small frame sizes, but at the same time, they suffer an increased PER. Given the transmission of 1024-Byte packets, the optimal rate adaptation should follow the envelope of the throughput vs SINR lines drawn in figures 4.6 and 4.7. For example, using IEEE 802.11b with a SINR at the receiver of 8 dB, the transmitter should use 5.5 Mbps CCK, but if the signal quality drops to 4 dB, the transmitter should change to 2 Mbps DSSS. Using IEEE 802.11a/g and with SINR = 10 dB, the transmitter should use the 18 Mbps modulation, but if the signal quality improves 8 dB, the performance could be optimized if 36 Mbps modulation is used instead. Note that the 9 Mbps modulation (BPSK, 3/4 coding) is never able to improve the performance of the 12 Mbps (QPSK, 1/2 coding), and therefore this modulation is never used, at least in an AWGN channel (used to draw the figures).

As stated in [252], an ideal link adaptation scheme should keep track of the SINR in reception, but this value is unknown for the transmitter, unless a kind of communication between transmitter and receiver is implemented. The authors of [163] presented the Receiver-Based Auto-Rate (RBAR) protocol, which proposes modifications in the IEEE 802.11 standard MAC in order to allow the exchange of information regarding channel conditions by means of RTS/CTS messages: first, the receiver estimates the wireless channel quality at the end of the RTS reception, then selects the appropriate transmission rate based on this estimate and feeds back to the transmitter using the CTS. A different approach is based on heuristics and estimations carried out locally in the transmitter side: the Auto-Rate Fallback (ARF) protocol [175] and its improved revision, the Adaptive ARF [194], keep track of a timing function and missed ACKs; in [104], the physical rate is adapted to the current link conditions as perceived by the transmitter through received signal strength measurements. In [184], a combination of these two approaches, missing ACKs and measured signal strength, is proposed: Hybrid Auto-Rate Fallback (HARF). In [73], a short survey on rate adaptation mechanisms is provided.

Most of the IEEE 802.11 MAC implementations have their own proprietary link adaptation schemes, but they are mainly based on heuristics inspired by the original ARF, even though it is well known that these schemes cannot react quickly when the wireless channel conditions fluctuate. The fact that these algorithms are not available to the public hampers the inclusion of rate adaptation in realistic analytical throughput models. For these reasons, the main objective of this section is to obtain a theoretical model for a generic rate adaptation algorithm, representative of real life implementations. To this end, in this section we study the behavior of a simple algorithm that counts successful/erroneous consecutive frame transmissions: after S consecutive correct transmissions, the rate is increased, and after E consecutive erroneous transmissions, the rate is decreased. This behavior is easy to model and simplifies the resulting formulation, while keeping our results representative of many rate adaptation implementations. The main difference with known implementations is the presence of a timer. In collision prone channels [61], missing ACKs can also be produced by collisions: even if channel conditions are good and the current modulation is optimal, the rate can be reduced due to collisions; decreasing the rate leads to higher transmission times thus rising collision probability. The use of timers in acknowledgment-based heuristics is intended to avoid rate reductions as a result of collisions. The absence of a timer in our model is justified by the fact that collisions are not taken into account in the rate selection.

The effect of collisions on throughput performance is already included in our multicell throughput model (cf. 4.3). Recall that the proposed algorithm is intended for modeling purposes and not for a practical implementation, although it would render a good performance under low collision probability.
The rate used to transmit any frame in a system with M modes can be modeled by the combination of the Markov chains shown in Fig. 4.11 a (main chain, MC) and b (secondary chain, SC). Both chains are actually coupled but are studied separately to simplify the resulting formulation. In MC, for any mode x, the transitions P_{mx}^{-} and P_{mx}^{+} between adjacent states represent the rates of change to mode $x-1$ or $x+1$, respectively. These rates are given by solving the lower level chains (SCs). There is a SC embedded in all mx states of MC. Within SC, after h consecutive transmission errors, state he is reached; the system reaches any ks state after k consecutive successfully transmitted frames. State (0) of mode x is used as a starting point and as a link state between adjacent modes: starting with mode x, state (0) of mode $x+1$ is reached after S consecutive transmissions; after E consecutive erroneous frames, the system switches to state (0) of mode $x-1$.

If P_{st} represents the probability of occurrence of state st, the previously mentioned P_{mx}^{-} and P_{mx}^{+} will be $P_{(E-1)e}PER_m$ and $P_{(S-1)s}(1-PER_m)$ respectively. Note that the $PER_m(SINR)$ is different for each modulation m given the SINR, as detailed in Section 3.2; and that computation of stationary probabilities is not affected by the fact that the Markov renewal process is actually embedded at slot boundaries. By solving SC, the probabilities of the different states can be written as a function of P_{1e} and P_{1s}, as follows:

$$P_{1e} = \left(P_0 + \sum_{i=1}^{S-1} P_{is}\right) PER_m \tag{4.18}$$

$$P_{1s} = \left(P_0 + \sum_{i=1}^{E-1} P_{ie}\right) (1 - PER_m) \tag{4.19}$$

$$P_{he} = P_{1e}PER_m^{h-1}; \qquad 1 < h < E \tag{4.20}$$

$$P_{ks} = P_{1s}(1 - PER_m)^{k-1}; \quad 1 < k < S \tag{4.21}$$

$$P_0 = P_{1s}(1 - PER_m)^S + P_{1e}PER_m^E \tag{4.22}$$

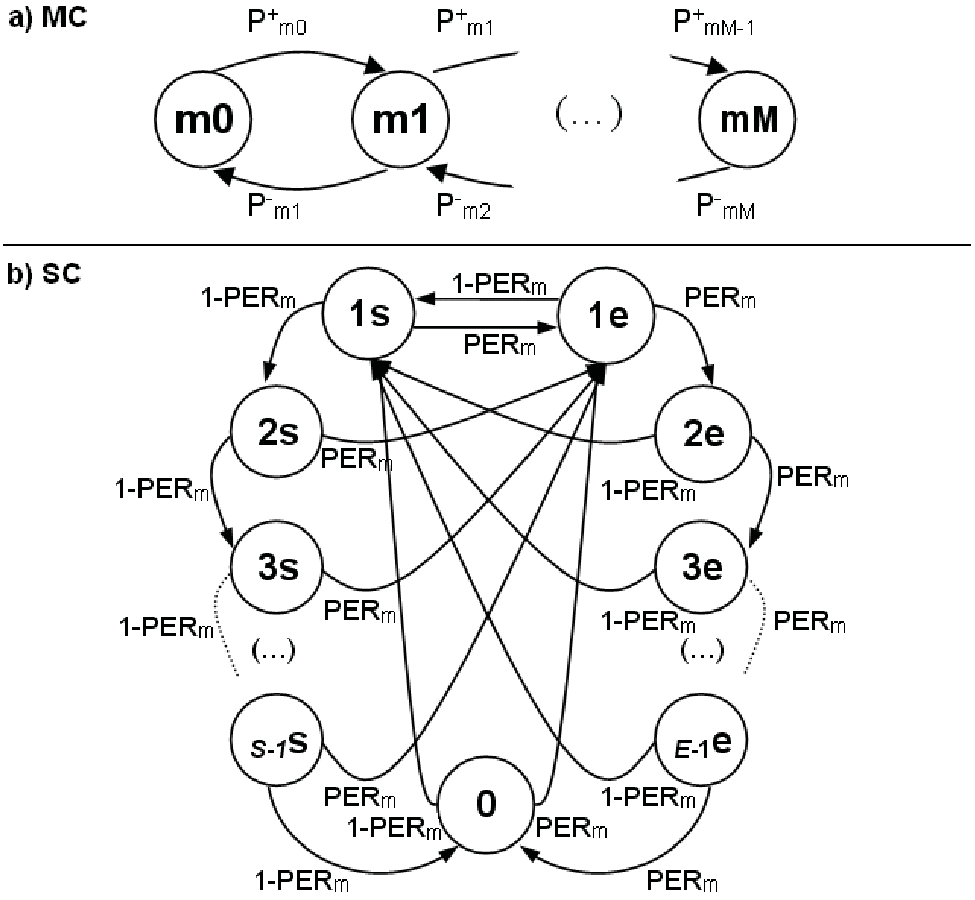

Figure 4.11: Markov chains modeling transitions between PHY modes

Note that P_0 actually depends on PER_{m-1} and PER_{m+1}, but in eq. 4.22 it appears as a function of PER_m. This error is due to the simplification introduced by the decoupling of the MC and SC chains. However, according to simulation results, this error is always smaller than 1%. In these simulations the rate adaptation algorithm described in this section is used to send 10^6 1000-Byte packets. Results obtained with the simulator are compared against the model in Fig. 4.12.

The sum of all probabilities (P_0, $P_{ks}\forall k < S$, and $P_{he}\forall h < E$) must be 1. From the relation between P_{1e} and P_{1s} depicted in 4.24 and derived from equations 4.19 to 4.22, we can leave all probabilities as a function of PER, S and E:

$$P_{1s} = P_{1e}\frac{1 - PER_m^E}{1 - (1 - PER_m)^S} \tag{4.23}$$

$$P_{1e} = \left(\frac{1 - PER_m^E}{1 - (1 - PER_m)^S}\sum_{i=0}^{S-1}(1 - PER_m)^i + \sum_{i=0}^{E-1} PER_m^i\right)^{-1} \tag{4.24}$$

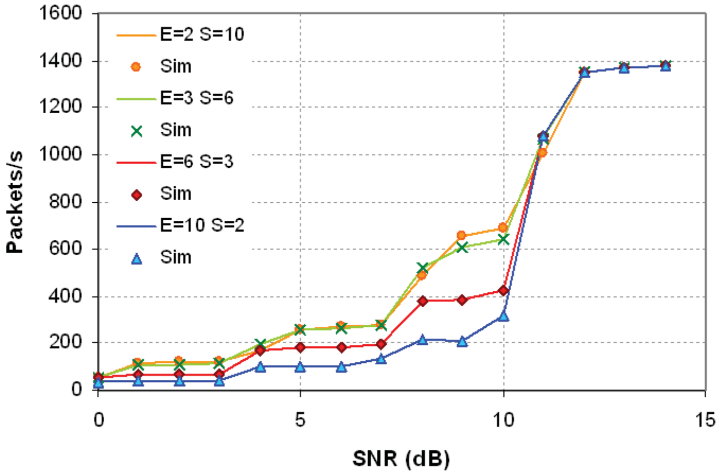

Figure 4.12: Rate adaptation model and simulations

Once we obtain the transition probabilities for all modes (P^-_{mx} and P^+_{mx}), the resolution of the chain MC is straightforward and allows us to compute the average PER and the average transmission rate as the weighted sum in eq. 4.27:

$$P_{m0} = \left[1 + \sum_{i=1}^{M}\left(\prod_{j=1}^{i}\frac{P^+_{m(j-1)}}{p^-_{mj}}\right)\right]^{-1} \tag{4.25}$$

$$P_{mx} = P_{m0}\prod_{j=1}^{x}\frac{P^+_{m(j-1)}}{p^-_{mj}}; \quad 0 < x \leq M \tag{4.26}$$

$$PER = \sum_{i=0}^{M} PER_{mi}P_{mi} \tag{4.27}$$

The above formulation provides the PER and rate of a given station. The next step is then to obtain the average PER and rate of the cell. Using these values in equations 4.11 to 4.13, as explained in Section 4.2, we can include the effect of rate adaptation into the saturation throughput model (eq. 4.17). As mentioned early, averaging PER and PHY rate in a cell is not required if we applied the formulation given in [210].

The best choice of E and S is settled after evaluating the model under different values of PER. Figure 4.13 is a sample of those tests: the three figures represent the number of packets successfully transmitted (normalized by the total number of transmitted packets) by a single IEEE 802.11b station operating at a given modulation and using different combinations of S and E. As expected, with high signal quality, the best performance is obtained with small values of S (and a bigger E). On the other hand, with a high loss rate, it is preferable to have a small E (and a bigger S). This trade-off is solved

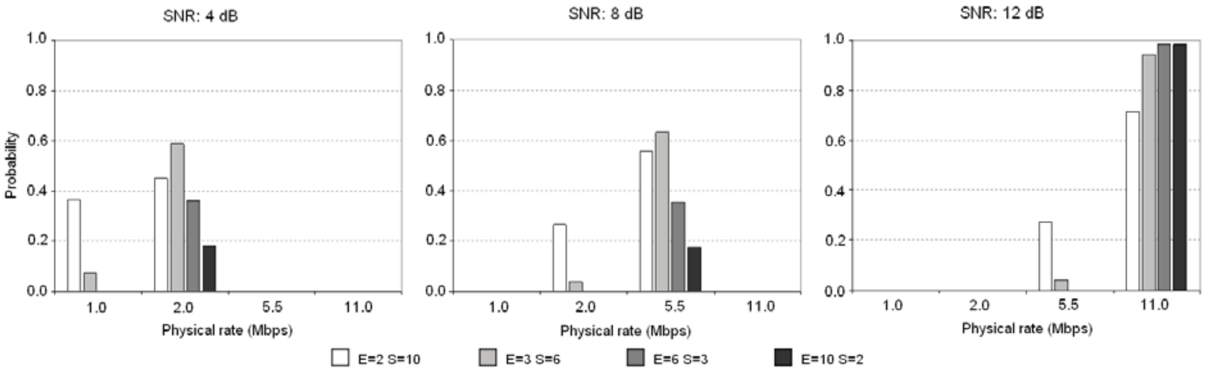

Figure 4.13: Histogram of successful transmitted packets for different SNR values and different (E, S) combinations

by selecting E=3 and S=6, which provides a good performance for a wide range of SNR values.

4.3.4 Practical evaluation

In order to evaluate the reliability of the rate adaptation algorithm modeled in the previous section, the goal of the first set of tests was the comparison of the expected performance of the algorithm with measurements obtained in a real scenario. The scenario consists of two laptops equipped with different NIC models configured in automatic rate selection. One CBR UDP connection is set to send 1500 Byte packets in order to maximize throughput (recall that packet size has no impact over the rate selection algorithm performance). Samples of throughput and SNR are taken while the distance between the laptops is slowly increased.
The models used are: 3com OfficeConnect (chipset ZyDAS ZD1211), Buffalo AirStation WLI-CB-g54 (chipset Broadcom 4306), and Allied Telesyn AT-WCL452 (chipset Intersil Prism 2.5). Prism's rate adaptation is the only algorithm whose basic principles are public. Prism cards implement a rate adaptation algorithm [10] similar to ARF: rate fallback is triggered by transmit retries; a timer function and successful transmissions increase rate.

As shown in Fig. 4.14, the mean throughput values obtained with practical measurements are lower than expected. This is not only due to the use of different rate adaptation algorithms. Indeed, taking measurements in a real testbed involves facing undesirable effects from environmental factors that negatively affect the measured throughput. While the analytical model only takes thermal noise into account, interference from other systems could have affected the testbed performance. Moreover, SNR measurements provided by the wireless device firmware can be inaccurate. However, the values provided by the model can be considered as an admissible upper bound, which most of the time is within one

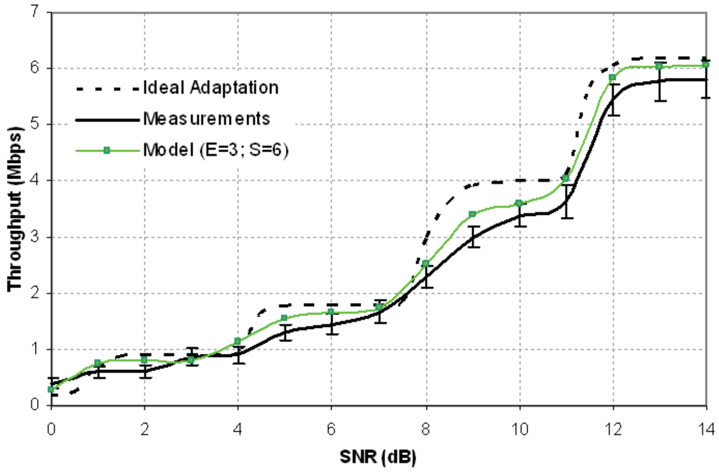

Figure 4.14: Rate adaptation model performance and real measurements

standard deviation away from the mean of measured samples. The dotted line represents the upper bound that could only be achieved with an ideal adaptation (i.e. instantaneous knowledge of receiver's channel conditions, cf. [104]). In order to validate the proposed model (eq. 4.17), including both adjacent and co-channel interference, another testbed has been built. Without loss of generality, we consider a scenario of three IEEE 802.11b overlapping cells (A, B and C) with two nodes each (AP and client), since it allows the analysis of the combined effect of adjacent and co-channel interference to be analysed, and hence remains applicable in larger scenarios. For each client, an upstream UDP source is set as in the previous tests. Channel 1 is set in cell A, C uses channel 11; the aggregate throughput of the three cells is measured for different channels of B (1 to 11). The first results obtained with our model using a fixed rate of 11 Mbps correspond to Fig. 4.15 line b; real measurements are depicted by line a. Once we include rate adaptation (line c), it is clearly shown that our approach closely models the measured results, even though the rate adaptation used in the model is slightly different from the algorithm actually implemented in the wireless devices.

4.4 Conclusions

This chapter presents a review of the research devoted to providing estimations of the capacity available in IEEE 802.11 WLANs. At first, the literature explored the limits of a single transmission. This evolved with the addition of transmission errors. Later, new models appeared that included the effects of collisions in multi-user scenarios. In spite of its assumptions, Bianchi's model is the most widely used model of this kind. This model has been revised by many researchers in order to include the effects of errors and to avoid some of the original assumptions. However, Bianchi's model was limited to a single-cell

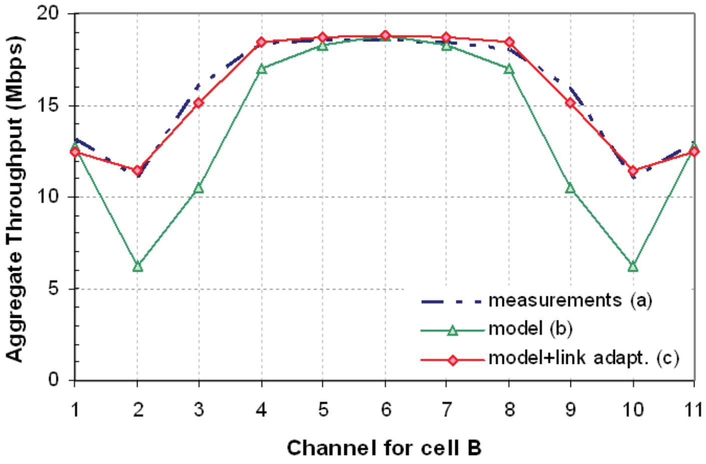

Figure 4.15: Aggregate throughput for cells A, B and C for different channel configurations in cell B

(single-channel) scenario. That is to say, it does not take into account the interaction between neighboring cells and the undesired effects of inter-cell interference. Therefore, new approaches were still needed to provide capacity estimations for large-scale WLAN networks (ESS).

For these reasons, we developed the multi-cell WLAN capacity model published in [130] as part of this thesis. We present a new method to obtain the maximum capacity of a multicell IEEE 802.11 network, considering the effects of co-channel and adjacent channel interference. An analysis of a simple method for the link adaptation is also included, which helps to provide more realistic results. Despite the fact that the rate adaptation algorithm presented is only valid for modeling purposes, its performance is shown to be representative of real-life devices, which implement proprietary algorithms.

In spite of the simplicity of the proposed algorithm and the considered assumptions, it has been demonstrated through simulations and real measurements that this algorithm provides a reasonably accurate estimation of the saturation throughput. Therefore, our model is a helpful tool for the qualitative evaluation of a multicell WLAN scenario. In this regard, the next chapter presents a practical application of the model: a frequency assignment scheme for IEEE 802.11 WLANs that uses the output of the capacity model as the target parameter to be optimized.

Chapter 5

Dynamic frequency assignments in WLANs

As the number of WLAN users grows, the need to perform efficient radio resource management strategies becomes essential due to the fact that most popular technologies, those based on IEEE 802.11 standards, use unlicensed frequency bands. A good channel assignment improves the network performance, producing benefits that are perceived by the users and also by the network administrators.

After an in-depth review of the related literature, in this chapter we present the approach developed during this thesis and published in [134, 138]: a new frequency management scheme for IEEE 802.11 WLANs in the 2.4 GHz ISM band that minimizes interference to increase the throughput available to client stations by adapting a weighted Degree of Saturation (DSATUR) algorithm for graph coloring. The algorithm tries to optimize the performance of the WLAN network by looking at the capacity estimations presented in Chapter 4. Recall that these estimations depend on both the co-channel and the adjacent channel interference, as studied in Chapter 3. In consequence, we make use of all available channels instead of the traditional nonoverlapping three. In this way, collisions as well as transmission errors are minimized, thus improving the network capacity and the user experience. Different architectures are discussed for the implementation of our approach, including the possibility to incorporate client stations into the management system.

5.1 Related work

The increasing demand for wireless communication brings the need to use new technologies. There are currently two main approaches to enhancing wireless capacity. The most common approach is to improve the radio interface. Technologies such as adaptive modulation coding, MIMO or OFDM are recent improvements that are able to increase the bitrate, but the ceiling on spectral efficiency will make further improvements difficult if not impossible. The only approach that can offer continuous improvement is the channel reuse according to a cellular pattern. Using smaller cells increases capacity, but it also makes radio resource management a more complex challenge.

In most of the literature, radio resource management in WLAN 802.11 networks is studied as part of the design of multicellular infrastructure-based WLANs. A good design is achieved if two basic requirements are met: full coverage of the required area and a capacity suitable for supporting the offered traffic is provided without service degradation as the number of users increases. Although there is an endless list of parameters to consider, the requirements mentioned above can be met using an exhaustive selection of AP locations and the proper set of channels and power levels. Many authors have contributed to solving this problem by developing different algorithms [292]. Rodrigues [260] and Hills [161] pioneered the application of these techniques to IEEE 802.11 WLANs. Since then, most of the work has been based on a thorough knowledge of the scenario: traffic demand in the different areas to cover, orography and obstacles, etc. An offline calculation is then made to evaluate certain parameters and configure the APs before they are deployed. In [197, 260] these optimization problems are described in terms of Integer Linear Programming (ILP), but other heuristics can be found in the literature, since the optimization problem is usually *NP*-hard. For example, in [199] authors try to minimize the effective channel utilization (portion of time that an AP finds the channel busy) at the bottleneck AP.

In a more realistic scenario, the possible interfering sources are unpredictable, and the number of APs, client nodes and the traffic demand vary randomly. Thus, the initial settings should be adapted dynamically to fit the new environment. Therefore, the frequency assignment problem requires a new mechanism capable of evolving according to new conditions and supporting the various WLANs in a given area. Planning is an effective solution for the deployment of a sole multicellular network: in [162] and [291] first fixed AP locations are set and subsequently, frequency channels are assigned. In the real world, however, network administrators know neither when nor where a new interfering node is going to appear.

These premises make the implementation of centralized schemes difficult, as discussed in [125]. Therefore, a distributed mechanism in which all nodes involved participate may be preferable. For example, in [211], each AP chooses its working channel by itself, minimizing the number of co-channel neighbors. Authors in [107] simplify this approach to allow each AP to set a new channel asynchronously. In [224] APs uncoordinately perform a “slow” frequency hopping in order to improve fairness among APs that will hence share the time spent in “bad” channels. However, this frequency hopping entails frequent reassignments and a consequent performance decrease. These kind of solutions provide local optimums; if cooperation is enabled among the APs, better assignments are obtained [259].

The algorithms for the Frequency Assignment Problem (FAP) can be classified according to their objective. When the purpose is to find a feasible solution that fits certain restrictions, we use the feasible frequency assignment problem (F-FAP) or maximum service FAP (Max-FAP) variants. This is the method most widely used in 802.11 networks; assignments that satisfy the 5-channel separation for adjacent cells are found [197,215,260]; [203] provides a Max-FAP solution, maximizing performance by minimizing co-channel interference. Optimizing the number of channels (MO-FAP: minimum order FAP) or the area of the frequency spectrum used (MS-FAP: minimum span FAP) does not make sense, since the entire ISM band is freely available. For these reasons, we argue that the most appropriate approach for the WLAN 802.11 distributed coordination function is the one that minimizes interference between cells, thereby improving the performance of each cell. This strategy has been used in [108], but this scheme is able to avoid interference by changing the AP placement. Moreover, the model used captures very roughly the interference and does not take into account the evident effect of utilization. The FAP formulation which aims to minimize interference is known as minimum interference FAP (MI-FAP), and is usually modeled using interference graphs $G = (V, E)$. The interference graph G consists of a set of vertices V, representing the WLAN cells in a region, interconnected by a set of edges E. The existence of an edge between two vertices indicates that the two cells are overlapping. There are various methods for solving these graphs [41, 42], many of which are designed to find optimal solutions. However, they all involve a large computational effort. The graph coloring heuristic methods stand out due to their simplicity and low computational overhead [109, 186]. Graph coloring tries to assign "colors" to the vertices of a graph so that neighboring vertices (i.e. those connected by an edge) have different "colors". The DSATUR algorithm [77], which was also used in [215, 259], provides a fast, simple heuristic. It establishes the order in which vertices are colored and the colors assigned. At each iteration, the vertex with the highest saturation degree (i.e. having the largest number of colored neighbors) is selected to be colored. If two or more vertices have the same saturation degree, the one with the highest ordinary degree (i.e. having the largest number of neighbors) is selected; if the draw persists, then a random selection is performed. Since the assignment may be computed on several different devices, all of which must obtain the same result, all nondeterministic steps must be replaced; e.g. in [215] ties are broken using the physical addresses. The color assigned to the selected vertex is the lowest channel not being used on any of the vertex's neighbors. The steps are summarized as follows:

1. Arrange the vertices by decreasing order of degrees.
2. Color a vertex of maximal degree with color 1.
3. Choose a vertex with a maximal saturation degree. If there is an equality, choose any vertex of maximal ordinary degree in the uncolored subgraph.
4. Color the chosen vertex with the least possible (lowest numbered) color.
5. If all the vertices are colored, stop. Otherwise, return to 3.

When there is a high density of nodes and edges (e.g. when the graph contains a clique of 4, as in Fig. 4.9, and we have just 3 colors), this algorithm is not useful since it tries to solve a problem that has no feasible solution. That is to say, like F-FAP, it is limited by the use of just three colors (i.e. three non-interfering channels) in 802.11b/g networks. To overcome this limitation, authors in [103] and [76] allow the use of "forbidden" colors by evaluating a certain cost. In [222], for example, the cost is limited to the number of clients in conflict (i.e. affected by interference from neighboring cells), but the degree of the interference is not evaluated. The idea of adding costs led us to the conclusion that partially overlapping channels could be used in IEEE 802.11 WLANs [134], once the effects of interference are studied. Later, in [258] authors also allowed partially overlapping channels by simply applying a straightforward table-based mapping. Even though in [179] the authors assume nonoverlapping channels, their formulation can be extended to the case of overlapping channels by taking into account the received power between interfering pairs of APs. Similarly, in the recent [112], use of overlapping frequencies is penalized depending on the channel distance and the distance between AP pairs (not taking into account utilization or transmitted power). In [46], a so-called overlapping channel interference factor is used to minimize inter-AP interference. However, according to their model, channel utilization has no impact on the interference computation, and adjacent channel interference is always preferable to co-channel interference.
In conclusion, a more accurate interference characterization was still needed to support the idea of using overlapping channels. In this regard, results in [129, 225] provided the key to open the whole ISM band for IEEE WLANs (see Section 3.3 for full details).

5.2 Proposed approach

As mentioned before, the main contribution of our approach lies in the fact that we make use of all the available spectrum, including partially overlapping channels. We first introduced this idea in [134] and later it was further developed and extended in [138].

Besides the interference-aware concept, we must provide a solution for a distributed and cooperative environment (i.e. each AP is capable of choosing the appropriate channel in cooperation with other nodes), and a centralized solution for the operator-centric approach of [243, 244]. In both cases, assignments must be dynamically provided in order to fit the frequency configuration into the varying nature of the radio environment. For these reasons, the proposed algorithm must be fast in providing new assignments and hence there is an inherent trade-off between the goodness of the channel set (the optimal solution may be expensive in terms of time and/or resources) and the level of adaptation to the rapidly changing environment.

5.2.1 Partially overlapping channels

Partially overlapped channels are a useful resort when the number of non-overlapping channels available is small. Thus, its applicability is mainly focused on the 2.4 GHz band, for both 802.11b's DSSS and 802.11g's OFDM; recall that although not all the channels defined in the 5 GHz band are non-overlapping, the number of non-overlapping channels in this band is large in contrast to that of the 2.4 GHz band (near 20 for 5 GHz and 3 or 4 for 2.4 GHz, depending on the regulatory domain).

A simple experiment can be carried out to illustrate why the use of partially overlapped channels could have a positive effect. The scenario consists of two totally overlapping WLAN cells, each composed of one 11b AP and one client station. Cell A reaches the saturation state (i.e. there is always a frame ready to be sent in the transmission buffer). Cell B's offered throughput is increased from 0 to saturation. Figure 5.1(a) shows the carried throughput in A, measured at the application layer for different channel settings: while A remains in channel 1, B moves from 1 to 6. In Figure 5.1(b), the same experiment is repeated after moving one of the cells 10 m away to an adjacent room. It is clear that co-channel interference can be worse than adjacent-channel interference, depending on the cell utilization and on the interfering signal level. For example, in the case that 4 or 5 channel separation is not possible due to a high density of WLANs, in the first scenario (Figure 5.1(a)), a network administrator would choose the same channel for both cells before setting partially overlapping channels. However, in the second scenario (Figure 5.1(b)) the received co-channel interference is still above P_{th} ($-76 < P_{th} < -80$ dBm), while a reduction of the interfering signal has led to an improved BER performance for adjacent-channel interference, to the point that a channel distance of 3 would be the best choice herein. For the same reason a channel distance of 1 ($c = 1$) performs better than $c = 2$ in the first scenario, but $c = 2$ outperforms $c = 1$ in the second. Thus, it is sometimes preferable to use partially overlapping channels.

Given that the power of a signal received at a distance of d can be computed as $P_{rx}[dB] = G(c) + P_{tx} - k\,10\,log10(d)$, where P_{tx} is the transmitted power, k a factor that depends on the environment ($k = 2$ for open space), and $G(c)$ a factor that captures the effect of different gain/loss elements (including tx/rx filters and antennas), isolating d:

$$d(c) = 10^{\frac{G(c)+P_{tx}-P_{rx}}{10k}} \tag{5.1}$$

we obtain d_{th} when $P_{rx} = P_{th}$, i.e. d_{th}, which is also known as the carrier sense range, is the minimum distance between two nodes required to prevent the carrier sense mechanism from reporting a channel occupied when the other station is transmitting. Note that $d(c)$ decreases with c, which means that the minimum distance required to avoid interference between two transmitters (due to CSMA/CA), is reduced by increasing the channel distance ($d_{th}(0) > d_{th}(1) > \ldots > d_{th}(c)$). Again we can say that it can be preferable to use partially overlapped channels before a channel that is already in use. Let us describe a new scenario to depict these statements. The scenario in Fig. 5.2

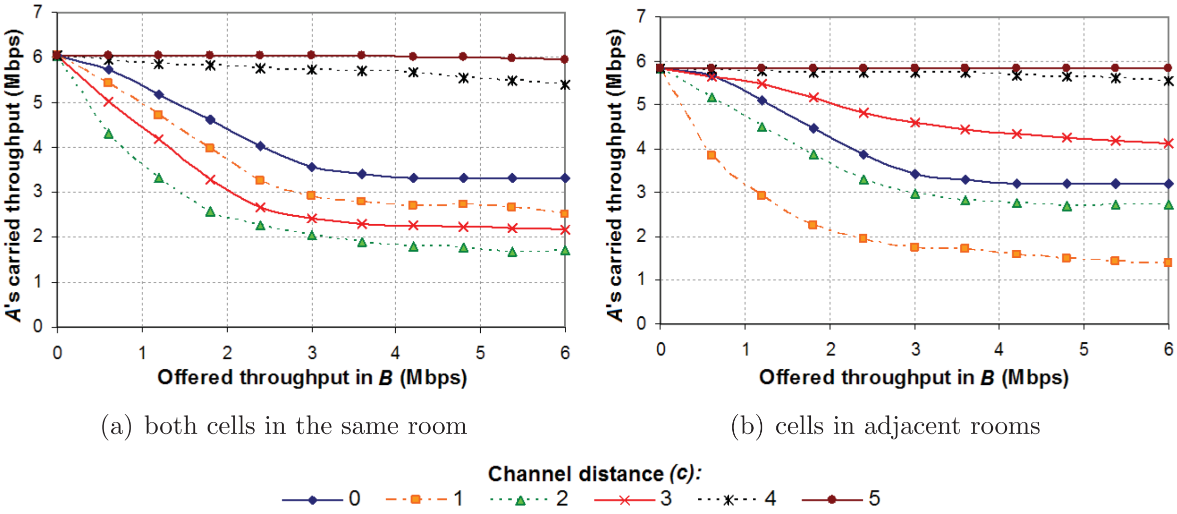

(a) both cells in the same room

(b) cells in adjacent rooms

Figure 5.1: Measured throughput in A with interference from different (partially overlapping) channels

contains four .11b tx/rx pairs. All transmitters are within carrier sense range of each other ($d_{xy} < d_{th}(0)$). Receivers are only within range of their corresponding transmitter. Hence, the best channel assignment for this scenario (with $max(d_{xy}) < d_{th}(0)$) will be that using channels 1, 5, 9 and 13. Observe that in this case, the use of partially overlapped channels avoids the exposed node problem (more on the exposed node problem in 3.3.5).

5.2.2 Algorithm

The two main requirements for our channel allocation algorithm are fast adaptation to changes and a low degree of complexity, thus allowing the algorithm to be run in low-featured devices (e.g. the AP itself). Therefore, our algorithm must work in a timely manner at a low computational cost. Moreover, the algorithm must be able to maximize the network capacity by making use of all the available channels. For these reasons we adapt the DSATUR with costs to fit the particularities of IEEE WLANs. This algorithm does not always provide the optimal solution, but it is fast and requires few resources. However, as mentioned above, if we use overlapping channels in interfering cells, performance degradation occurs. We therefore made some modifications to the algorithm in order to reflect the effects of interference. In our approach, the cost is based on the capacity estimations presented in Section 4.3.

The saturation degree retains its definition from the original DSATUR: number of colored neighbors. The concept of ordinary degree is modified: the new ordinary degree evaluates the interference of a node seen from its neighbors, including its utilization, as

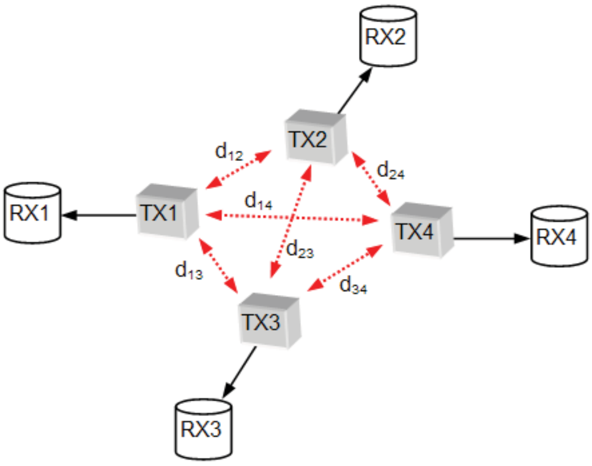

Figure 5.2: Avoiding exposed node problem in 2.4 GHz WLAN

explained in Section 3.3.4. In the simplest case study, let us assume that each vertex of the graph corresponds to an AP, that is, V represents the set of APs; an edge between two vertices denotes that the two APs are within reach of each other. To compute an interference-based cost, we then use a matrix C whose elements $c[i][j]$ are defined as follows:

1. $c[i][j] \quad \forall i, j \in V$ represent an average of the signal level received in AP j from the cell of AP i.
2. $c[i][j] = 0$ means that in graph $G = (V, E)$, the edge $(i, j) \notin E$.
3. $c[i][i] = 0 \quad \forall i \in V$.

The sum of the elements of the ith row of C gives the total interference caused by node i and is used in our algorithm as the ordinary degree. The interference level received by a node could also be used likewise, obtaining similar results. Also note that all uncontrolled APs in the graph (not participating in the process) must be considered as colored nodes from the beginning. The idea is represented in algorithm 1.

First of all, the nodes are arranged by decreasing order of degree; the degree of a node is decided by two criteria: 1. saturation degree and 2. new ordinary degree. Then all APs in V are colored in order. We define $f(v)$ as the channel assigned to the AP v. The function $bestChannel()$ assigns a channel to v such that the capacity of the colored cells is maximized. In this way, we do not define any “blocked colors” and so there is always a feasible solution, in contrast to other versions of DSATUR. Capacity estimations are based on a simplification of equation 4.17. The function $bestChannel()$ can be written as:

$$bestChannel(v, G) = argmax_{f \in F} \sum_{i \in V : f(i) \neq 0} S_i \tag{5.2}$$

Algorithm 1 channel assignment for IEEE 802.11 WLANs

```
Require: $(G, C, F)$
Ensure: $f(v) \neq 0 \forall v \in V$
  while $!(f(v) \neq 0 \forall v \in V)$ do
    select $v$ with $f(v) = 0$ and MAX Sat. degree
    if there is a tie then
      select $v$ with $f(v) = 0$ and MAX Ord. degree
    end if
    $f(v) = bestChannel(v, G)$
    update Sat. degree $\forall v \in V$
  end while
```

where F is the set of available channels and S_i is obtained from eq. 4.17, but only taking APs and their cross-interference into account in this case. Here, the values u_j are known in advance, so the specific heuristic explained in Section 4.3.2 need not be applied. Note that the sum of S_i does not actually represent network capacity unless all nodes are driven to saturation.

It turns out that the choice of the first channel assigned has considerable impact on the quality of the assignment. Since no generally good rule could be identified as to which channel should be assigned first, the algorithm is actually run $|F|$ times, each time starting with a different channel, and finally the best assignment is chosen.

5.2.2.1 Client-driven enhancement

Two distant APs (out of decodable range of each other) can belong to interfering cells if there is a client station within transmission range of both APs. In an AP-only approach, the graph will not reflect the actual interference. However, it should be recalled that the original cost function obtained with equation 4.17 also includes interference as seen by clients. The forthcoming IEEE 802.11k amendment will allow APs to request from their clients a list of in-range APs and other radio measurements that better describe the actual interference. Moreover, the optimization metric will be more accurate once we include the information provided by the clients. The next section presents more details on this issue.

5.2.3 Implementation issues

As mentioned above, the channel allocation mechanism is implemented in two ways: centralized and distributed. Due to the great diversity of WLAN networks that may coexist in a given area, it may be difficult to coordinate all nodes in order to obtain optimal frequency assignments [45]. This makes centralized management solutions difficult to implement (cf. [125]). A distributed system of some sort would therefore be appropriate, but this would require additional intelligence for the participating devices. This requirement is supported by current technology trends, which are expected to keep producing devices with improved features, i.e. new APs will have faster CPUs, greater storage/memory capacity, etc. Nevertheless, it must be remembered that current off-the-shelf APs still have limited features. Therefore, we take several steps to reduce the complexity of our algorithms so that the computational cost will be affordable enough for execution in commercial APs. For example, although capacity estimations are accurate enough for our needs, some errors are introduced due to the assumptions made in the model. Consequently, there is no need to fully implement the complexity introduced by Bianchi's formulation in the function $S_{th}()$ (see Section 4.3.2). In fact, $S_{th}()$ returns tabulated values from a matrix whose coordinates represent the average PER and the number of contenders. As a result, the computational cost is reduced without further affecting accuracy. This makes the algorithm run faster, which allows several executions for a given scenario, introducing and evaluating small variations in the solution set. Thus, the algorithm is able to produce better assignments. These variations are deterministic in a distributed implementation and random in a centralized architecture, as with simulated annealing optimizations.

In contrast, centralized management systems rely on a central entity capable of running complicated processes without hindering the normal operation of the managed devices. Centralized solutions simplify many issues that are still immature in the distributed solutions (e.g. security, communication among members, etc.). The next subsections will focus on the implementation of the distributed and centralized mechanism. We then discuss various client-driven techniques that are valid for both centralized and distributed management.

Either way, the central manager (in a centralized scheme), or the APs (in a distributed solution) need to build the interference graph in order to run the algorithm. In both cases, whoever is in charge of performing the channel assignment must be able to write out a text representation of the interference graph. This graph will be read and interpreted by the piece of software implementing the allocation algorithm. The text representation for interference graphs used in our implementation is an extension of the Discrete Mathematics and Theoretical Computer Science (DIMACS) format[1]. Figure 5.3 shows an example of graph and its text representation. The line beginning with n specifies the number of nodes in the graph, u the utilization of each AP (in %), f the initial channel with which each AP is configured. The e lines describe the edges of the graph: $e\ 0\ 1 - 60 - 65$ denotes that there is interference between node *0* and *1*; *1* receives *0*'s *beacons* at -60 dBm, and *0* receives -65 dBm from *1*. A utilization or interference level is zero when it is not known.

The information required to build the graph is shared among the APs or sent to a central entity by means of some sort of signaling mechanism. This issue is discussed in the following sub-sections and detailed in Chapter 7.

[1] http://dimacs.rutgers.edu

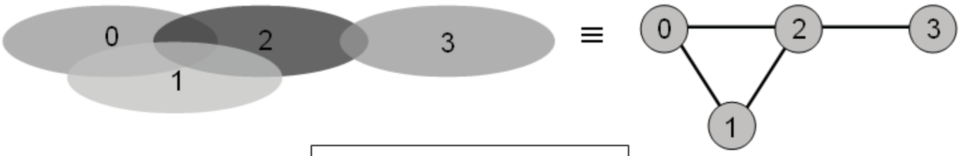

```
n    4
u   75  75   75   0
f    1   6   11  11
e    0   1  -56 -58
e    0   2  -45 -49
e    1   2  -38 -41
e    2   3    0 -83
```

Figure 5.3: Text representation of an interference graph

5.2.3.1 Distributed architecture

In order to run the simplified version of the capacity estimation algorithm explained above, an interference graph must be built in all APs. To do so, all APs must send information to all other members, including their own utilization and the signal power they receive from neighboring cells (without any filter).
The problem of adding signaling capacity to a group of elements that cooperate to utilize a shared resource better was studied in [127] and is discussed in Section 7.3.2. We argue that, although most WLAN networks are infrastructure-based (use of APs), from an inter-network point of view, their behavior could be compared to that of ad hoc networks. Following this idea, our scenario consists of a varying number of APs that can appear and disappear unpredictably in a region. Although its immediate application is limited by the boundaries of a single administrative domain, cooperation among independent domains ruled by different owners is also desirable when they share resources, as with the unlicensed ISM frequency bands used in IEEE 802.11 WLANs. In these cases, cooperation allows better resource utilization than greedy behavior by the competing domains. However, inter-domain cooperative resource management is difficult to establish due to the potential for malpractice and other security issues.

In [133], we implemented the simplest solution: all APs are connected to a common wired broadcast domain, but if the infrastructure is expensive to use or there is no common infrastructure at all, this solution is not suitable. The possible lack of a common wired distribution system makes us “think wireless”, so we propose using APs with more than one wireless interface: one is devoted to keep working as an AP, and the others can be used to join a wireless mesh network for inter-AP communication. It is a good idea to use a different technology/frequency band for the extra interfaces in order to avoid new interference (e.g. IEEE 802.11g for AP service and 802.11a for signaling). In addition to signaling, inter-AP communication could also be used to build a wireless backbone, as proposed in [299], which would allow the additional cost introduced by the extra interface to be redeemed.

Moreover, the lack of wires makes it easier to rapidly satisfy communication needs in temporary locations such as emergency sites, sporting events, etc.

To allow new custom software to be developed and tested in commercial APs, these devices should incorporate an open-source development environment. The main problems found during implementation were due to differences between driver developers, which are often inevitable due to differences between manufacturers in the way the physical layer is accessed. For example, signal levels provided by different firmwares are usually represented in different scales with non-standard units. This makes interference calculation difficult and causes erroneous results. In addition, not all models can perform frequency scans when acting as an AP or even report any type of signal measurement. Chapter 7 provides more arguments and expands this discussion.

5.2.3.2 Centralized architecture

The dynamic channel assignment strategy detailed in this chapter was included in the project UMTS-Assisted Mesh Network (UAMN) [243, 244] (see Section 7.2.3). This approach consists of a centralized architecture, where all APs are used to provide Internet access to Wi-Fi clients following the traditional infrastructure-based mode. However, APs are inter-connected through a wireless mesh network, which is in turn connected to a cellular operator through a UMTS or GPRS interface. On the operator side there is a server devoted to centralizing different management services, including frequency assignments.

In this centralized architecture, the required signaling is limited to the reports sent from the APs to the central unit. These reports are contained in UDP datagrams that are sent either periodically or triggered by changes in the interference graph. As in the previous section, the required information that is sent by the APs to the central unit include the current utilization of the cell, the list of APs in range and the received signal power.

If the central unit detects that a channel switch could improve the capacity of the network by a preset threshold, it sends the new channel assignment to each AP.

5.2.3.3 Client-driven

Next-generation networks are distributing core functionalities towards the boundaries of the operator. Like the authors of [223], we argue that terminals may hold the key to improving resource management. In our case, interaction with client devices is justified by the need to build a complete interference graph. Ideally, client devices should send information, including the interference received from both neighboring clients and APs within range, to the management system. Recall that IEEE 802.11k-enabled devices will be able to send these interference reports to their APs. With this information, we can use the original formulation instead of the simplified version used previously. The work in [243, 244] proposed an alternative to 802.11k, which was under development at that time: UMTS-Wi-Fi hybrid terminals could use the UMTS interface to exchange signaling with a central unit in order to assist in various key management functions.

Client-network cooperation is also required to smooth AP channel switching. We noticed that a sudden channel change can be traumatic for active clients (c.f. [64]). After the AP performs a channel change, active clients do not receive layer 2 ACKs, so they use up their maximum number of trials and activate the RTS/CTS mechanism, assuming there is a hidden node problem. If no answer is received, the clients start an active scan until the AP is reached. The transmissions are then resumed without repeating the association and authentication process. The total duration of connectivity loss is between 200 and 450 ms, depending on the card manufacturer, the traffic load, and the channel distance. IEEE 802.11h amendment establishes new mechanisms that an AP can use to announce a channel switch to all of its clients before it is actually performed (more details in Section 7.1.1). This synchronization will prevent harmful connectivity losses.

5.2.4 Evaluation

The RRM strategy detailed in this chapter logically improves the network performance, since it effectively reduces the harmful effects of interference. As detailed in the following sections, the improvement becomes more evident as the number of active users and their load increase.

First, we evaluate the losses of the chosen heuristic, comparing the output of our algorithm against the optimal solution. The results support our decision to sacrifice the optimal solution, obtained at a high cost, by a "good" solution that is rapidly obtained at a lower cost.
Next, the improvements provided by our approach are quantified by means of simulations and measurements in a small testbed.

5.2.4.1 Algorithm performance

First of all, in order to evaluate the goodness of the algorithm, a *C* program was made to generate random graphs. By solving a large number of these graphs, we verify that if a feasible solution is obtained with original DSATUR (i.e. up to three channels are needed to avoid interferences), it will also be obtained using our modification with exactly the same cost. When DSATUR cannot provide a feasible solution (e.g. the graph is not 3-colorable), the use of all possible channels always results in a less interfered system. As argued in Sections 3.3.4 and 5.2.1 and shown in the following paragraphs, this effectively translates into an improved capacity.

In addition, the assignments provided by the proposed algorithm are compared with the optimal solution (obtained using a *branch-and-cut* method) in order to measure its effectiveness and efficiency. Paper [45] uses statistics gathered in different US cities to provide useful hints for building WLAN interference graphs. Those statistics present extreme values, e.g. graphs with more than 8000 nodes, and nodes with a degree higher than 80 (i.e. an AP with more than 80 neighbors). However, the graphs are not connected, and even though the maximum connected component can have a size of hundreds of APs, the average connected graph has 10 to 16 APs. Moreover, more than half of the nodes have only two neighbors or less, which gives a low graph density. The graph density, represented by D ($0 \leq D \leq 1$), is defined in eq. 5.3 as the number of edges in the graph under study ($|E|$) divided by the number of edges in a complete graph with $|V|$ vertices. A density $D = 1$ means that the graph is a complete graph where every pair of distinct nodes is connected by an edge.

$$D = \frac{2 \cdot |E|}{|V| \cdot (|V| - 1)} \tag{5.3}$$

Therefore, to represent real-life dense scenarios, we build random connected graphs with $10 \leq |V| \leq 30$ and $3/(|V| - 1) \leq D \leq 8/(|V| - 1)$, so that the APs have 3 to 8 neighbors on average. The results show that our algorithm obtains the optimal solution in almost 40% of the tests (see Fig. 5.4); in 90% of the tests the error is below 10%. Here the error is defined as follows: if a is the capacity estimation obtained with the optimal set and b is the metric obtained with our algorithm, the error is $(a - b)/a$. Logically, $a \geq b$. An error of 0 means that the algorithm provided the optimal solution.

On the other hand, our algorithm is able to provide a solution set in 1 - 3 s running on a commercial AP[2] while the optimal set takes from 1-2 minutes for small graphs, to more than 10 hours for graphs with 20 or more nodes, when the *branch-and-cut* algorithm runs on a server station[3].

5.2.4.2 Practical measurements

The improvements obtained using a channel assignment strategy aimed at increasing the available capacity per station was previously proved in [133, 134] by means of practical measurements and implementing an early version of our mechanism: in a small testbed with 4 APs we obtain from 10 to 16% of increased capacity.

Measurements taken in the first testbed are devoted to showing the significance of cell utilization, as against frequency management schemes that only take received signal strength into account. The configuration of the testbed, represented in Figure 5.5(a), is the simplest topology where interferences cannot be avoided: a clique of four nodes. It is not solvable with F-FAP and three colors (i.e. three non-overlapping channels). The testbed consists of four APs and one client associated to each AP. Coverage areas

[2] CPU MIPS 400 MHz, RAM 64 MB

[3] 2xCPU Intel Xeon DC 2.4 GHz, RAM 2 GB

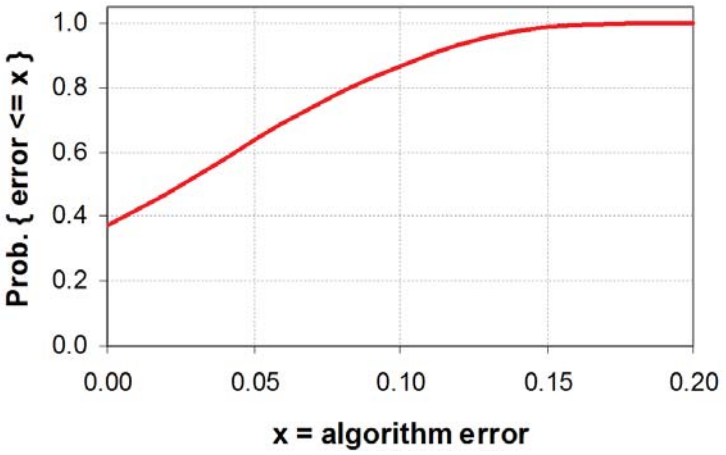

Figure 5.4: Cumulative distribution of the error of the proposed algorithm, against the optimal solution

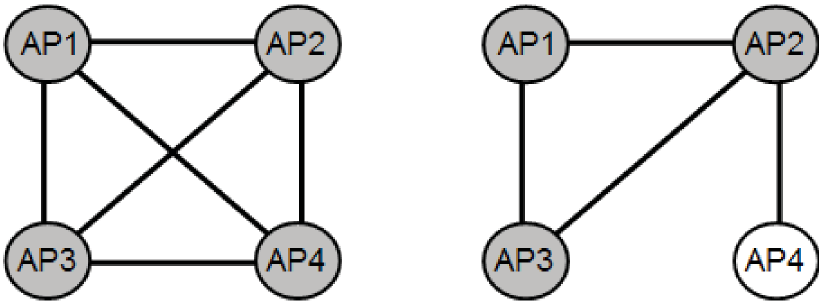

(a) 1^{st} frequency assignment testbed (b) 2^{nd} frequency assignment testbed

Figure 5.5: Interference graphs for practical measurements

of the eight nodes are overlapped, i.e. every node "sees" all the others. The results and configurations of the tested scenarios are summarized in Table 5.1. In the initial situation, one of the APs transmits at a high power level and generates low traffic, while the rest of the APs radiate low power but are affected by a high traffic load. If only received power levels are taken into account, the best assignment would minimize the interference caused by $AP1$ (1^{st} assignment). Considering not only power levels, but also traffic load, the assignment provided by our algorithm improves the global throughput (2^{nd} assignment). Then, the behavior of the nodes is altered in the following way: $AP1$'s transmitted power is reduced and its traffic is increased; the rest of nodes lower their traffic but their transmitted power remains unchanged. The previous assignment may not be the best in the new situation (1^{st} assignment of scenario 2), so a new assignment is computed, by which a better performance is achieved (2^{nd} assignment scenario 2).

A second set of measurements demonstrates the advantages of dynamically tunning

Scenario 1			**1^{st} assignment**			**2^{nd} assignment**		
AP	**Mbps**	**Power**	**Node**	**Chanel**	**Carried traffic**	**AP**	**Chanel**	**Carried traffic**
1	0.1	High	1	1	**10.18 Mbps**	1	1	**11.79 Mbps**
2	4.0	Middle	2	6		2	1	
3	4.0	Middle	3	11		3	11	
4	4.0	Middle	4	6		4	6	

Scenario 2			**1^{st} assignment**			**2^{nd} assignment**		
AP	**Mbps**	**Power**	**Node**	**Chanel**	**Carried traffic**	**AP**	**Chanel**	**Carried traffic**
1	4.0	Low	1	1	**7.66 Mbps**	1	1	**8.46 Mbps**
2	2.5	Middle	2	1		2	6	
3	2.5	Middle	3	11		3	11	
4	2.5	Middle	4	6		4	6	

Table 5.1: Configuration and results of 1st frequency assignment testbed

channel frequencies in a scenario where not all the APs participate in the frequency management. The testbed consists of three IEEE 802.11b APs implementing our frequency management scheme following a distributed architecture (c.f. [133]). Two of them are commercial devices (Linksys WRT54G, see section 1.2), whereas the other AP is built with a Pentium II 350MHz PC with a WLAN NIC based on Prism 2.5 chipset allowing master mode under Linux. A fourth distant AP completes the scenario. This fourth AP emulates an AP which is not under our control. The graph representation of the testbed is shown in Figure 5.5(b). $AP1$ and $AP3$ are commercial devices, $AP2$ the PC and $AP4$ the "uncontrolled" AP.

Initially, only $AP4$ is up and operates on channel 11, following a randomly variable traffic pattern. Then $AP1$ is turned on. It does not detect any interfering AP and decides to work on channel 1. Then, a client is associated and starts to generate traffic causing the utilization of $AP1$ to increase up to 75%[4]. $AP2$ initiates and after detecting $AP1$, it decides to use channel 6. A client is then associated to $AP2$ causing a 75% utilization. Next, $AP3$ is started; it detects 3 APs using the three non-overlapping channels. It decides to set channel 11, since it receives less interference from $AP4$. A client is connected to $AP3$ and the total throughput is measured. Throughput of cell 4 is "not known" to us, and hence not used in our calculation. Results and configuration are summarized in Table 5.2. After these measurements, obtained following a distributed but non-cooperative frequency selection, $AP1$, $AP2$ and $AP3$ collect the required information to build the complete graph and run the algorithm. After the new assignment, throughput is increased (see Table 5.2).

[4] 75% utilization is equivalent to 4.3 Mbps of UDP CBR traffic

Scenario 1			**1^{st} assignment**			**2^{nd} assignment**		
AP	**Mbps**	**Power**	**Node**	**Chanel**	**Carried traffic**	**AP**	**Chanel**	**Carried traffic**
1	4.30	Middle	1	1	**4.30 Mbps**	1	1	**4.30 Mbps**
2	4.30	Middle	2	6	**4.05 Mbps**	2	11	**4.13 Mbps**
3	4.30	Middle	3	11	**3.52 Mbps**	3	6	**3.98 Mbps**
4	?	Low	4	11	**? Mbps**	4	11	**? Mbps**

Table 5.2: Configuration and results of 2nd frequency assignment testbed

5.2.4.3 Simulations

Here we present a more complete study in which the benefits of our approach are evaluated through a vast number of independent simulations, providing small confidence intervals, which are therefore not shown in the figures. The simulator used, SIPTE11, implements all the details specified in the IEEE 802.11 standard, including the effect of interference in a multicellular environment, as detailed in Section 3.4.2.
The simulated scenario consisted of a 250x250 m^2 area where 20 APs and a varying number of randomly placed client stations. Using a semi-open urban propagation model (see Section 3.1), with all stations (APs and STAs) transmitting 15 dBm with omnidirectional antennas, the transmission range varies from 25 to 30 m due to the randomness introduced by the propagation model. Thirteen channels were made available, only 3 of which were non-overlapping. The load offered by each station was also set at random.

In the first set of simulations, we set the ratio of saturated stations (always having 1500 byte frames in the transmission queue) to 10%. In this case, all four solutions tested (random allocations, the proposed algorithm, 3-coloring and 4-coloring scheme) were able to carry the offered traffic. However, the harmful effects of interference became more noticeable as the number of active users increased (see Fig. 5.6). The proposed solution minimized both adjacent channel interference (thus improving PER) and co-channel interference (reducing collision probability). With random allocations, the PER was very high, but on the other hand, collision probability was better than with 3- or 4-coloring because the number of cells using the same channel was low.

If all stations try to obtain the maximum possible throughput (100% saturation), a frequency management strategy aimed at reducing interference increases carried traffic by about 18%, as shown in Fig. 5.7; i.e. a higher cost means a higher capacity (cf. Section 4.3). Note that even a random allocation improves the network performance obtained with the traditional 3-coloring.

With low loads, the presence of radio resource management is not noticeable; hence, network performance is the same for all channel allocation strategies. As the network load increases, the need for good radio resource management becomes critical. In these scenarios, the performance of our proposal stands out. Although our channel assignment mechanism supports the coordination of APs from different domains (cf. [127]), in the

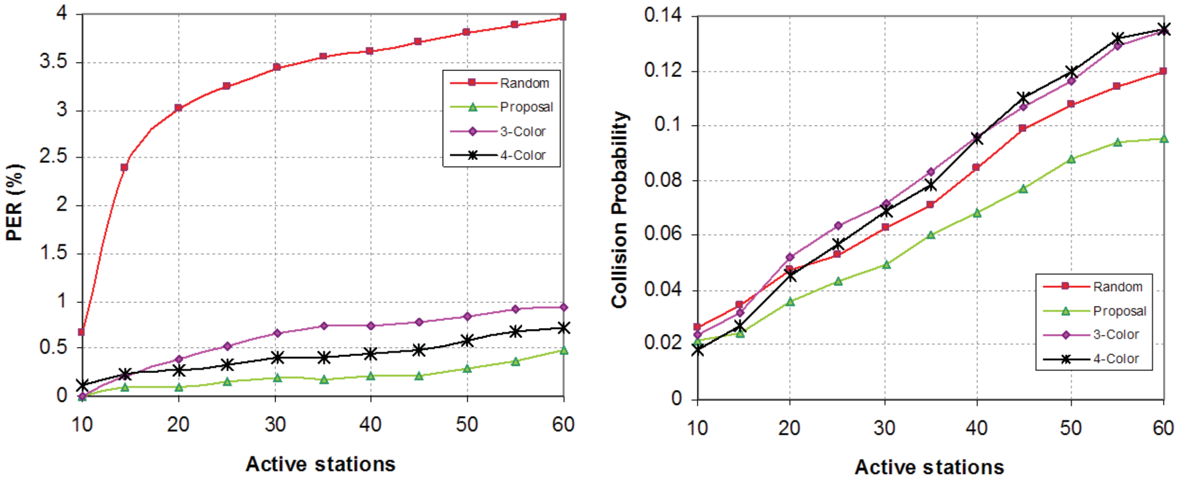

(a) PER vs Number of active users (b) Collision probability vs number of active users

Figure 5.6: Co-channel and adjacent channel interference mitigation (10% sat.)

real world, the presence of interfering cells from outside of the managed set of APs is inevitable. From the managed network's point of view, these sources are treated as a fixed-channel node in the interference graph. Logically, performance improves as the number of managed APs increases, but as shown in Fig. 5.8, under low load conditions (low ratio of saturated users, e.g. "*10% SAT*"), controlling 40-50% of the APs is enough to optimize network performance. With higher loads, at least 70% of the APs should participate in frequency management.

Finally, client-driven enhancements are verified in the last set of simulations. Figure 5.9 shows the carried traffic with different a number of saturated stations (10, 50 and 100%), using both the basic (AP-based) and client-driven algorithms. These improvements, which reach a maximum of 6%, are most visible when the network load is high.

5.3 Conclusions

In this chapter, we have introduced a new mechanism that solves the frequency assignment problem in IEEE 802.11 WLANs. We use all available channels, rather than just the traditional non-overlapping three, in order to effectively reduce both co-channel and adjacent channel interference.

We justify the use of the entire available spectrum, including partially overlapping channels, by studying the effects of adjacent channel interference through practical measurements in a small testbed. Based on the capacity estimations derived in the previous

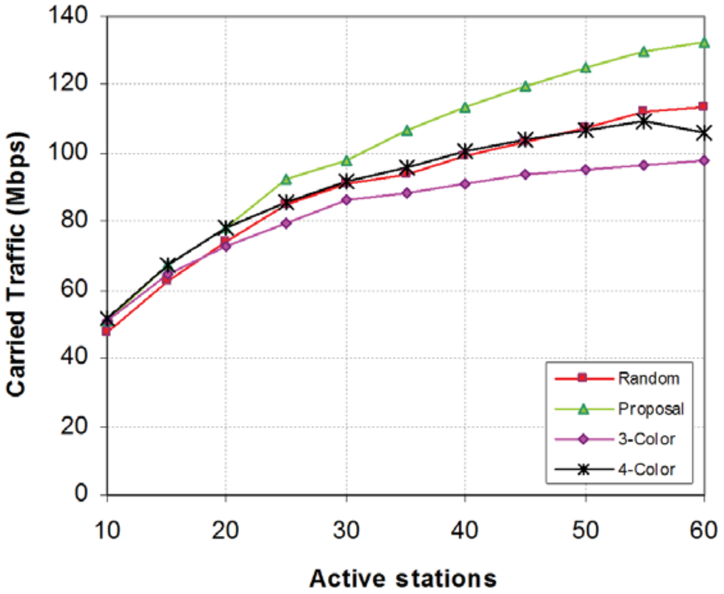

Figure 5.7: Carried traffic in saturation conditions

chapter (section 4.3), we build a cost function that can accurately evaluate different channel assignments. Recall that these estimations take into account the effects of co-channel and adjacent-channel interference adequately for a multi-cell WLAN scenario. We demonstrate that optimizing our cost function effectively reduces both types of interference, which logically increases the capacity of the whole network. This improvement becomes more noticeable as the load is intensified.

The benefits of good frequency management can be further amplified by integrating the role of client devices into the system thanks to the valuable information they can provide to build a complete interference graph. Both approaches-basic (AP-based) and client-driven enhancement-can be implemented following either centralized or distributed paradigms. In fact, all possible architectures have been implemented as proof of concept [134, 244].

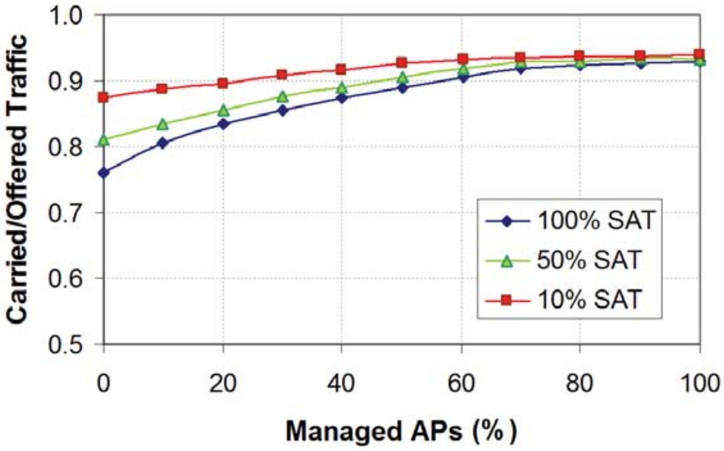

Figure 5.8: Carried traffic with varying number of managed APs and saturation levels

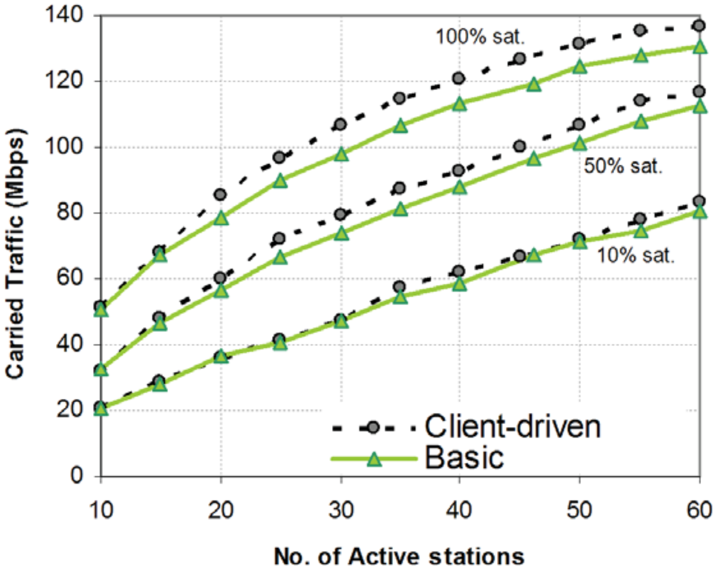

Figure 5.9: Client-driven enhancements on carried throughput

Chapter 6

Load Balancing in WLANs

Various studies have shown that WLAN user service demands are highly dynamic and vary for each type of network, although certain long-term patterns can be identified [56, 57]. These studies conclude that users tend to be concentrated both temporally and spatially, creating highly congested areas known as Hot Spots. For example, the volume of wireless Internet access in a convention center increases during coffee-breaks at the APs located close to the area where the coffee-break takes place. Therefore, the load is unevenly distributed across a small number of APs in the WLAN. Moreover, although mobility is increasing as users get into the habit of using wireless access, the mobility pattern can still be considered quasi-static [158] in the sense that users tend to remain in the same location for long periods.

This situation is compounded by the fact that the associations are determined by the client devices on the basis of signal level measurements, which means that users are generally associated with the closest AP. In other words, although a Hot Spot is served by several APs, most of the users will be connected through the AP that provides the strongest signal. Note that the maximum traffic carried by an IEEE 802.11 AP is limited and that this limit decreases as the number of active users increases due to the higher collision probability.

Following these considerations, we argue that by implementing different load balancing policies it will be possible to improve network efficiency and increase network ability to satisfy QoS requirements. First of all, it is essential that we define the concept of *load*. The next section provides an in-depth study of load metrics, concluding with the selection of the most complete definition: the Available Admission Capacity (AAC). This concept, published in [126,131], captures the effects of increasing the number of active users, varying the offered traffic and the channel quality. Furthermore, the AAC can be measured in real-time on actual commercial devices, as detailed below. Once the load is determined, two different load balancing techniques are proposed. First a client-driven approach published in [128], and in a second approach, load balancing is achieved by means of cell breathing techniques as detailed in [137].

6.1 Load Metrics

Load balancing in overlapping areas has traditionally been used in circuit-switched cellular networks. Since each user in these types of networks represents an identical utilization of available resources, load balancing could be applied by using call level information, i.e. a load balancing scheme will try to balance the number of active calls between neighboring base stations (BS).

Nevertheless, call level information is not sufficient for modeling the actual load carried by a BS in current wireless packet networks, given that users may have different traffic profiles. This assertion is valid for IEEE 802.11 WLANs and is corroborated by different empirical studies [56, 57, 158], which state that a new metric based on packet level information is required.

6.1.1 Number of active users

As shown in Fig. 4.5, the number of active users still provides valuable information in networks that use CSMA-based access: more collisions occur as the number of active users increases, which leads to decreased performance. In [122, 124], the number of associated STAs is announced by the AP. The number of competing stations, n, can be calculated from any station by using the formulation given in [71]:

$$n = 1 + \frac{log(1-p)}{log(1-\tau)} \tag{6.1}$$

where p is the probability that a packet collides and τ the probability that a STA transmits a packet in a randomly chosen slot time. Both parameters were previously defined in section 4.2. Each STA is able to derive this value if p (which is the same for all members of the same cell) is known. Recall that τ can be obtained as a function of p (see eq. 4.8). However, this parameter can only be used to estimate the saturation throughput of a cell and does not provide information about the actual load.

6.1.2 Packet loss

Different load metrics based on packet level information have been proposed. The authors of [105] used the number of retransmission attempts needed to successfully transmit a single packet, which can be derived if all hidden STAs are known. The same concept, but following a different formulation, was also used in [72] to derive the *Gross Load* (*GL*). The PER of any STA i may be known to the AP from the SNR (see section 3.2). Then the probability that STA i transmits a frame after j unsuccessful attempts is given by:

$$P_i^{succ}(j) = (1 - PER_i)PER_i^j \tag{6.2}$$

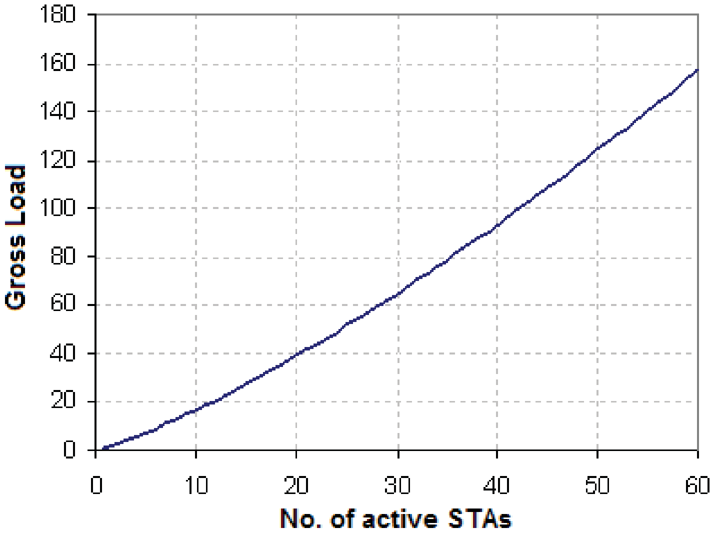

Figure 6.1: Gross Load (GL) vs. number of saturated users

According to these authors, the load contributed by STA i, g_i represents the number of slots needed for a successful transmission from i. Finally, the total load GL is the sum of all individual loads:

$$g_i = 1 + \frac{PER_i}{1 - PER_i} \tag{6.3}$$

$$GL = \sum_{i=1}^{n} g_i = n + \sum_{i=1}^{n} \frac{PER_i}{1 - PER_i} \tag{6.4}$$

Figure 6.1 shows the evolution of the GL as the number of saturated users is increased. Transmission errors are not considered and thus $PER_i = p\ \forall i \in N$. Reference [72] also suggests using the packet loss (PL) estimation as a new load metric. More precisely, they evaluate the cost, in terms of packet loss, of adding a new STA. Even though its computation is not trivial, it can be obtained from GL. In this way we can simplify the formulation assuming that all STAs have the same PER ($PER_{eq} = 1 - N/GL$). The average PL in a cell with M available slots and n STAs ($M = n + k$) is:

$$PL(N,k) = \frac{1}{n} \sum_{i=1}^{n} i(1 - PER_{eq})^{n-i} PER_{eq}^{i} \binom{n+k}{n-i} \tag{6.5}$$

In [123], besides the number of associated STAs, the authors include the effect of packet loss in client-driven association decisions. STAs infer the value of PER from RSSI local measurements following the empirical equation 6.6 (δ denotes the effect of channel distance and received power from interference source). Then the load of an AP is $(1 - PER)/n$.

$$PER = 20000 \left(\frac{7\, RSSI\, \delta}{8} + 94 \right)^{-3.8531} \tag{6.6}$$

6.1.3 Carried traffic

Traffic (in bytes/s) was used as a load metric in [55] and [48], where the absolute load is computed as the measured carried throughput divided by the saturation throughput; in [55], if this value exceeds a given threshold, the cell is considered to be overloaded. In [48], ILP is used to minimize the congestion of the most congested AP. In [280] and [167], the authors use relative load values by means of Jain's fairness index [168]. An AP is overloaded or underutilized depending on its neighbors' carried traffic. However, fairness in CSMA access entails the so-called "performance anomaly" [160] in the presence of multi-rate stations. We should bear in mind that the IEEE 802.11 standards define several modulations with different physical bit rates (e.g. 1, 2, 5.5 and 11 Mbps for 802.11b); in this case an AP could be congested when carrying traffic of 1 Mbps if there are associated stations transmitting at the slowest bit rate. On the other hand, the same AP could also be considered under-utilized with a load of 3 Mbps if all of its clients use faster modulations. Therefore, carried traffic is not a valid representation of the load on an AP in a multi-rate scenario.

Instead, in [281] authors propose the Multi-rate Performance Index, with which transmissions at different bitrates are taken into account. In [60] the measure of busy time is proposed as the representative load metric (see next section). More precisely, the network congestion level is estimated by using channel occupation time and by monitoring the occupation of the AP's buffer queue. Note that the AP queue size does not represent the overall network congestion level in the case of asymmetric traffic load since it does not depend on the uplink traffic, which also contributes to the resource saturation. However, downlink traffic generally overrides uplink traffic in many Internet applications.

The concept of *fairness* is also used in [190] as part of their strategy to balance load in an IEEE 802.11 WLAN. In this case, the authors seek a fair share of available bandwidth instead of a balanced load among APs. They try to optimize the function $\Omega(A)$ that depends logarithmically on the throughput of the n STAs:

$$\Omega(A) = \sum_{i=1}^{m} ln\left(\theta_i(A)\right) \tag{6.7}$$

where $\theta_i(A)$ represents the saturation throughput for STA i in the solution set A.

6.1.4 Load information in IEEE 802.11 standards

The IEEE 802.11k group (Radio Resource Measurement) recently released a new amendment intended to improve the provision of data from the physical and medium access

layers by defining a series of measurement requests and reports that can be used in the upper layers to carry different radio resource management mechanisms. There is more information regarding the IEEE 802.11k in section 7.1.2.
The current IEEE 802.11 standard [24] and the new 11k define a set of load metrics that are either broadcast by APs or measured directly by client stations:

Channel Load Report: any station accepting a *Channel Load Request* shall respond with a *Channel Load Report* that is defined as the proportion of the time during which either the physical carrier sense, the virtual carrier sense (Network Allocation Vector or NAV) or both indicate that the channel is busy. This measurement is similar, although not identical, to the CCA report added by the IEEE 802.11h.

Beacon Frames: these management frames are extended with three new elements that provide information about the load of an AP.

- *BSS Average Access Delay*: average medium access delay (MAD) for any transmitted frame measured from the time the frame is ready for transmission (i.e., begins CSMA/CA access) until the actual frame transmission start time.
- *BSS AC Access Delay*: in QAPs (QoS enabled APs), average medium access delay for each of the indicated Access Categories defined by the IEEE 802.11e (Best Effort, Background, Video and Audio).
- *BSS Load*: contains information on the current STA population and traffic levels in the BSS. Includes the following fields:
  - *Station Count*: the number of stations currently associated with the AP.
  - *Channel Utilization*: the percentage of time that the AP senses the medium is busy, which is indicated by either the physical or NAV mechanism.
  - *Available Admission Capacity (AAC)*: the remaining amount of medium time available via explicit admission control.

6.1.4.1 Service Time and Medium Access Delay

Service Time is usually defined as the time interval from the moment at which a frame is at the head of its MAC queue ready for transmission until it is successfully received at the destination. This definition includes lost time caused by collisions or transmission errors. The value can be calculated by using analytical models [85, 91, 257]. Nevertheless, the simplest definition can be derived from Bianchi's formulation [70] for saturation conditions: Service Time is the interval between two consecutive, successful transmissions performed by the same station in a cell with n active nodes:

$$T_{serv} = N \frac{E_s}{P_{tr} P_s} \tag{6.8}$$

where E_s is the average length of a renewal interval, defined as the time between two consecutive transmissions or the time between two consecutive backoff decrements, taking into account collisions, transmission errors, idle time and data transmission time. P_{tr} is the conditional probability that at least one transmission occurs in a randomly chosen slot time and P_s the conditional probability that this transmission is successful. See [70] and section 4.2 for further details. From eq. 6.8 it is clear that T_{serv} is proportional to the number of active users, but note that E_s and the product $P_{tr}P_s$ also depend on n. The dependency on n can be seen in Figure 6.3(a), where the effects of packet size are also noticeable, as expected. The figure is also useful for comparing analytical values -solid lines- with simulations -markers- (refer to section 3.4.2 for more details on the simulation environment).
The average service time is equal for all stations in the same cell due to the long-term fairness provided by the CSMA/CA access scheme, so any given station can estimate the average service time from local measurements.

Similar to the above parameter, 11k defines Medium Access Delay (MAD) as the average time a frame is held in the transmission buffer queue. MAD is measured from the time at which the DCF frame is ready for transmission (i.e. CSMA/CA access begins) until the actual frame transmission start time. If a frame needs to be retransmitted (because it has not been acknowledged) the value of MAD is calculated by averaging the waiting times of the r retransmission attempts (see Fig. 6.2):

$$MAD = \frac{\sum_{i=0}^{r} MAD_i}{r+1} \tag{6.9}$$

The conceptual differences between the two parameters are clearly shown in Fig. 6.2: the first parameter comprises the time needed to successfully send a frame, and MAD is the average time a frame is held in transmission buffers due to contention. Figure 6.3(b) shows the difference between Service Time and Medium Access Delay as the number of users is increased. The figure was obtained from the analytical model based on stations in saturation state (eq. 6.8) using 48 Mbps modulation and a packet size of 1024 Bytes. MAD was also derived from eq. 6.8 with the average number of retransmissions (eq. 6.3).

As stated in [282], these parameters capture the effects of multi-rate STAs and the number of active users, and are therefore useful in gauging the effective load of a cell. However, they can lead to misguided conclusions, since some scenarios with very different number of users may result in a similar load, according to these metrics. For example, as shown in Fig. 6.4, the MAD value in a cell with two active STAs only differs from the case with twenty when the offered traffic approaches saturation levels. Furthermore, in a cell with only one or two active STAs, MAD and service time measurements do not change noticeably as the offered traffic is increased.

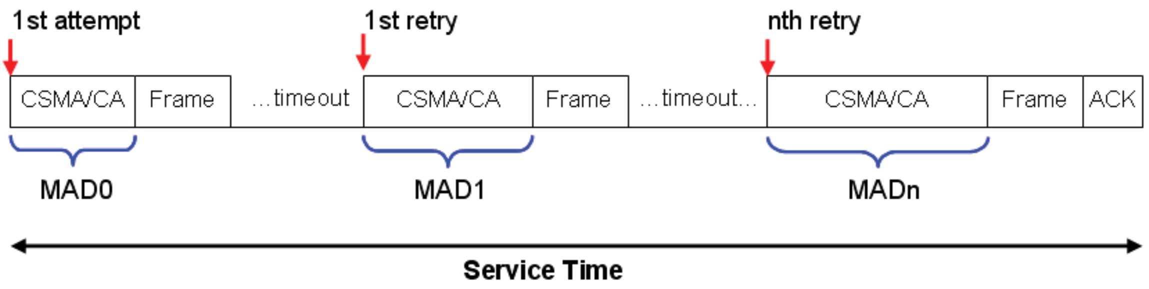

Figure 6.2: Service Time (T_{serv}) and Medium Access Delay (MAD)

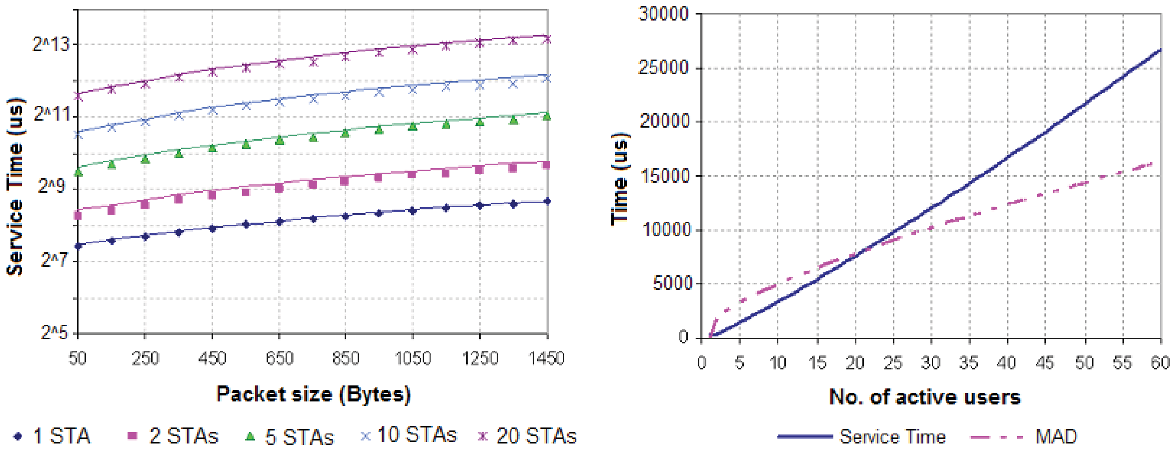

(a) T_{serv} for different number of STAs and packet sizes

(b) MAD, T_{serv} vs. number of saturated STAs

Figure 6.3: Service Time (T_{serv}) and Medium Access Delay (MAD) with increasing number of active STAs

6.1.4.2 Channel busy time

Another metric for the load in an IEEE 802.11 WLAN cell is the channel busy time. Channel busy time is the percentage of time during which the medium is sensed as busy, either due to a successful transmission or due to a collision. It provides more valuable information on the load when STAs transmit at different rates.

The medium is under-utilized when there are a small number of competing stations due to the backoff mechanism used in the IEEE 802.11 CSMA/CA implementation (refer to section 2.3 for more information). On the other hand, a larger number of competing stations will decrease the probability of finding an idle slot. Busy Time can be modeled using the formulation defined in section 4.2 for saturation conditions, as in the previous

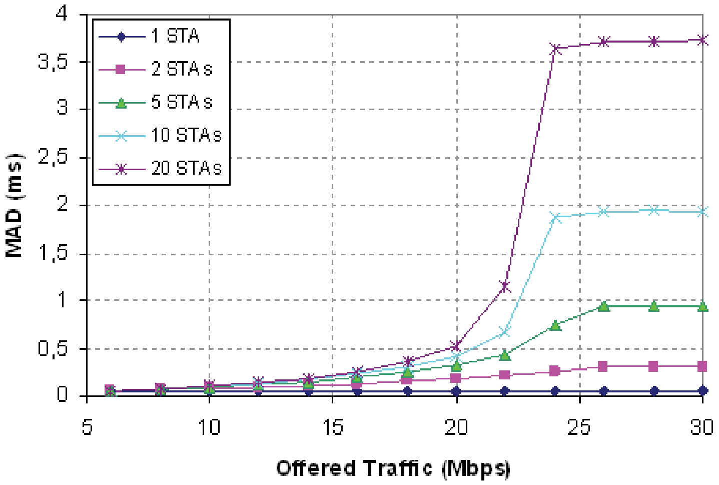

Figure 6.4: MAD with increasing offered traffic in a IEEE 802.11g cell

section:

$$T_{busy} = 1 - \frac{(1 - P_{tr})\sigma}{E_s} \tag{6.10}$$

E_s and P_{tr} have already been defined; σ is the slot time unit defined by IEEE 802.11 standards (20 μs for .11b and 9 μs for .11a/g).
Figure 6.5 illustrates these efficiency issues and shows how busy time is affected by the number of active users, the average frame size and the offered traffic. The figure is obtained from simulations using IEEE 802.11g standard values at 48 Mbps in a single-cell scenario. Figure 6.5 shows the evolution of this load metric as a function of different parameters; Fig. 6.5(a) is obtained through analytical models in saturation conditions and Fig. 6.5(b) is obtained from simulations, fixing a packet size of 1024 Bytes. In both figures IEEE 802.11g and 48 Mbps are used.

Although channel busy time provides a good representation of the cell load even in a multi-rate scenario, it is not a valuable metric in the presence of greedy applications, that is, applications that use as much banddwidth as possible (e.g. bulk FTP transfers). For example, a channel busy time of 85% is achieved with a single greedy station (Fig. 6.5(a)), but also with 20 users, each of which is loading the cell with 1.1 Mbps (Fig. 6.5(b)). However, a new station will get much more bandwidth if it only has to compete with one user than if it has to share the medium with 20 other stations.

A slightly different metric is used in [63]. It provides the same information although its definition differs from busy time: the load contributed by a STA i is the time that takes for an AP to provide STA i one unit of traffic. Thus, the load induced by STA i, L_i is computed $L_i = 1/r_i$ (r_i being the physical rate at which STA i is associated). The total load is then the sum of all L_i using the AP. In [110, 111], the same concept is used, but the authors include the effect of co-channel interference.

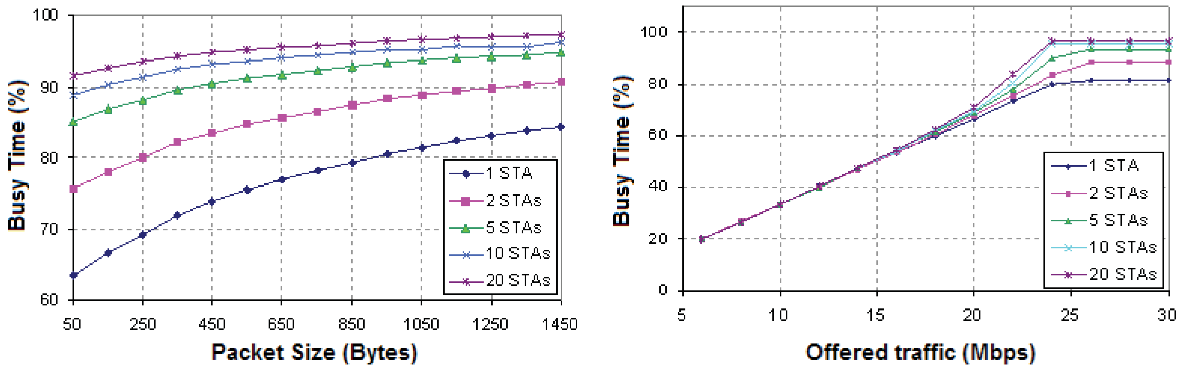

(a) Busy Time for different number of STAs and packet sizes

(b) Busy Time vs. offered traffic

Figure 6.5: Busy Time variations with number of users, offered load and packet size

6.1.4.3 Available Admission Capacity

The IEEE 802.11e standard defines Available Admission Capacity (AAC) as the remaining amount of medium time available in units of 32 μs, although it does not specify how it should be calculated. We propose a new load metric based on a more precise definition of Available Capacity. We expand the current definition of AAC to be used in non QoS 802.11 WLANs: AAC is the proportion of time a new station can take up if it is associated with the AP at a given physical rate. This concept has recently been used in [195] under the name of *Estimated aVailable bAndwidth (EVA)*.
This new metric provides a vision of cell load that takes into account the effect of multi-rate stations, the presence of greedy users, the average frame size and the number of active users. We provide a more detailed explanation in the following sections.

6.1.5 Available Admission Capacity

In this section, the computation of AAC is explained. As defined above, AAC represents the proportion of time a new station can take up if it is associated with the AP at a given physical rate. From this value, it is easy to predict the saturation throughput that that STA would obtain.
We argue that AAC values can be used to describe the load of an IEEE 802.11 cell, since this parameter depends on the number of active users, their offered traffic, the modulations used and the channel conditions.

Straightforward measurements can be used to derive the throughput that is currently devoted to a station with ongoing data transmissions. However, more detailed study is required to determine how to predict the throughput that a new station will obtain before it actually starts to transmit, or the maximum bandwidth that can be allocated to a station if it increases its offered traffic in a multi-rate cell. There are strategies that obtain those estimations by sending small amounts of probe traffic periodically (e.g. see [264]), but they are usually based on client-server applications. In a previous work [131], we proposed an algorithm that can accurately calculate the throughput available to a new station in a multi-rate WLAN environment. The process is totally transparent to the user and does not produce extra load. Amendments to this approach were added in [126], in which a comprehensive evaluation was also provided. Our definition of the AAC has been used as load metric in our work [128], [137], and also in other unrelated proposals by different researchers (e.g. [235] and [147]).

The information required to carry out the estimation of the AAC are easily available at the AP. This way, the AP is able to compute its own AAC by itself.

6.1.5.1 Definitions

This algorithm is based on the assumption that the IEEE 802.11 MAC maintains fairness in terms of access probability independently of the rate and bandwidth requirements of each station. It also takes into account the inherent "performance anomaly" [160] in multi-rate CSMA/CA networks. Based on these statements, we define T_{cycle} as the average time required to send one frame from each of the competing stations:

$$t_i = DIFS + T_{data}(i) + SIFS + T_{ACK} \tag{6.11}$$

$$T_{cycle} = max(TBO_i) + \sum_{i=1}^{n} t_i \tag{6.12}$$

Note that t_i is defined for a basic CSMA/CA access; if RTS/CTS handshake is used, 2·SIFS, T_{RTS} and T_{CTS} must be added to t_i. The duration of the data frame is T_{data}. TBO_i represents the average time during which a station i waits for the backoff timer to expire before attempting to transmit. Note that under saturation conditions, backoff slots are shared. Therefore, the idle time spent on backoff within a T_{cycle} is equivalent to the largest average backoff in the cell. Without saturation, stations are not "synchronized", and the probability that two stations share their backoff periods (or a part of them) is lower. Consequently, without saturation, our approximation introduces some error. In the next section we show that the error introduced is small.
TBO_i depends on the number of previous transmission attempts. The average value of the backoff interval after j consecutive transmissions is given by:

$$T_{BO}(j) = \begin{cases} \frac{2^j(CW_{min}+1)-1}{2} T_{slot} & 0 \leq j \leq 6 \\ \frac{CW_{max}}{2} T_{slot} & j \geq 6 \end{cases} \tag{6.13}$$

We consider that the number of packet retransmissions required to successfully transmit a single packet is a geometrically distributed random variable. If P_i is defined as the probability that a frame sent by i has to be retransmitted (including the effects of collisions and channel errors), the average backoff interval for station i is:

$$TBO_i = \sum_{j=0}^{\infty} T_{BO}(j)(1 - P_i)P_i^j \tag{6.14}$$

Note that from eq. 6.14 we have to compute an infinite series. However, its operation can be stopped when the required precision is met. For example, for typical P_i values ranging from 0.1 to 0.2 and a required precision of 1 μs, 6 to 10 iterations are enough. T_{data} can be further decomposed into[1]:

$$T_{data}(i) = T_{preamble} + \frac{8\,(H_{MAC} + MSDU_i)}{r_i} \tag{6.15}$$

where r_i is the physical bit rate at which node i sends data frames with a payload of MSDU bytes. The remaining values vary according to the standard and the modulation used. We define the overhead produced in layers 1 and 2 of a node i as:

$$OH_i = \frac{r_i\, t_i}{8\, MSDU_i} \tag{6.16}$$

In order to calculate the actual TO (traffic offered) it is necessary to consider this overhead and the possible retransmissions given the traffic offered by the upper layers (TO_{app}):

$$TO_i = \frac{OH_i}{1 - P_i} TO_{app}(i) \tag{6.17}$$

6.1.5.2 Algorithm

Resource distribution in WLAN provides a Max-min fairness in which small flows receive the volume they demand and larger flows share the remaining capacity equally. In other words, all stations whose traffic (TO_i) is equal to or smaller than the bandwidth that should be allocated under saturation conditions will be able to carry all of their offered traffic (TO_{app}). Saturated stations will share their corresponding bandwidth and the excess time that is not used by the non-saturated nodes. The process of finding the carried throughput for any station in a cell with N stations is represented in algorithm 2.

S_j is the throughput that station j would obtain in saturation conditions. The parameter δ_i is defined as the proportion of the maximum throughput that can be achieved by station i which is actually used: $\delta_i = TO_i/S_i$. Before the first execution, stations are ordered according to increasing value of the parameter δ. Stations with $\delta_i \leq 1$ use

[1] case of DSSS-CCK modulations, for more details on OFDM see [252].

Algorithm 2 AAC estimation

```
Require: $(T_{cycle}, \delta_i, r_i, TO_i \ \forall i \in N)$
Ensure: $L_i, S_i \geq 0 \ \forall i \in N$
  $T'_{cycle} \leftarrow T_{cycle}$
  OrderIncr($N$, $\delta$)
  for all $i \in N$ do
    if $\delta_i \leq 1$ then
      $L_i \leftarrow TO_i/r_i$
      $T_{exc} \leftarrow (1 - \delta_i)/r_i$
      $T'_{cycle} \leftarrow T'_{cycle} - 1/r_i$
      for all $i \in N$ and $\delta_i > 1$ do
        $L_i \leftarrow L_i(1 + T_{exc}/T'_{cycle})$
        $S_i \leftarrow L_i r_i$
        $\delta_i \leftarrow TO_i/S_i$
      end for
    end if
  end for
```

fewer resources (or an equal number) than those that would be allocated under saturation conditions, and their TO_{app} will therefore be carried. The proportion of time that is not used by non-saturated stations (T_{exc}) will be divided fairly between the remaining stations according to a new time T'_{cycle} in which the stations that have already been served are not considered. Note that greedy applications are modeled with $\delta_j > 1$ regardless of the value of the S_j parameter. The TO of a new station is not known in advance and the maximum throughput it can potentially achieve is therefore calculated in saturation. The value of L_i represents the individual load contributed by node i defined as the proportion of T_{cycle} used by the ith node.

By using this formulation it is possible to perform the capacity estimation in real time, assuming that the required statistics are updated regularly by the firmware/driver of the wireless interface. Consequently, this estimation can handle varying traffic demands and varying channel conditions, since it also takes into account the effect of collisions and errors produced by noise and interference.

6.1.5.3 Evaluation

In the following paragraphs, the throughput and AAC estimations provided by the algorithm described above are evaluated by means of analytical models, simulations and practical measurements. Results led to the conclusion that the assumptions made in our models can be held without affecting the accuracy.

6.1.5.3.1 Analytical models The well-known Bianchi model [70] allows an accurate evaluation of the saturation throughput of IEEE 802.11 DCF networks under the assumption of ideal channel conditions, and considering unlimited retransmissions, by employing a Markov chain (see section 4.2).

As mentioned above, in equations 6.11 and 6.12 we are introducing errors since we consider that the backoff slots are always shared among the stations. This is true under saturation conditions, that is, when all stations have always one or more frames ready for transmission. In this situation, all stations decrement their backoff timer at the same time and the one whose timer expires first gains the access; the remaining stations will decrement the remaining backoff time during the next idle period. Figure 6.6(a) shows the results provided by our algorithm (dotted lines) and results obtained with Bianchi's model (markers). Note that Bianchi's model imposes the following constraints for the scenario: no packet errors, no hidden nodes, and all stations use the same modulation. Throughput values are obtained for the four modulations used in IEEE 802.11b. The MSDU is 1500 Bytes. The average relative error is below 1%.

Chatzmisios et. al. introduced the effects of packet errors within Bianchi's model [90]. In Fig. 6.6(b) we show the results provided by our algorithm (dotted lines) and results obtained with Chatzimisios' model (markers) for different values of the PER.
The graph is obtained with 10 stations and frame size of 1500 Bytes. Recall that P_i is the probability that a frame is not successfully transmitted due to either errors or collisions. In this way, the relationship between P_i - equations 6.14, 6.17 - and PER is:

$$P_i = PER + P_c - PER\, P_c \tag{6.18}$$

where P_c is the collision probability. In this case, the relative error is also below 1%.

6.1.5.3.2 Simulations As shown in the previous sub-section, backoff slots are shared among saturated stations, but in a more general case, the presence of unsaturated stations avoids such a "synchronization". Two stations will share their backoff slots (or a part of them) only if a packet arrives at both transmission queues at the same time, or during a third station's transmission, or during each other's backoff period (given that the remaining time is greater than DIFS). The probability of this event is low for low offered loads and is increased as we approach saturation and the number of users is increased. That is to say, the error introduced in our approach is larger with low loaded cells.

The previous evaluation based on analytical models was limited in many ways: it can only be done under saturation conditions, no multi-rate environment and no hidden stations. Conversely, simulations allow a more comprehensive evaluation. Refer to Section 3.4.2 for more details on the simulation environment. The scenario consists of one IEEE 802.11g AP serving various clients. The physical rate used for transmissions depends on the distance between the STA and its AP. Figure 6.7 shows the correlation between rate and distance for the mandatory set and some optional 802.11ag modulations.

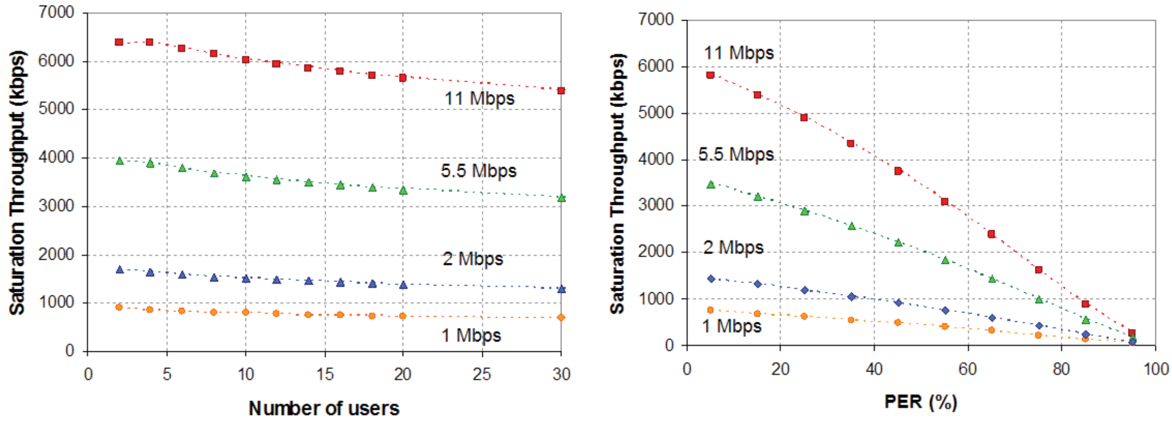

(a) Saturation throughput with increasing STAs

(b) Saturation throughput with errors

Figure 6.6: Throughput estimations with AAC algorithm (dotted lines) and saturation models (markers)

The radii are obtained by using the propagation model for semi-open offices detailed in Section 3.1.

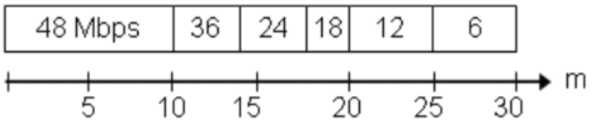

Figure 6.7: AP's coverage radius for different physical rates

The first set of simulations is intended to demonstrate that the "performance anomaly" of CSMA/CA-based WLANs [160] is effectively captured by our model. The scenario consists of a single IEEE 802.11g cell with an AP and three stations (A, B and C): A uses a physical rate of 48 Mbps to carry the requirements of a greedy application (FTP), B offers traffic described by a 7 Mbps UDP CBR source while transmitting at 24 Mbps, and C increases its traffic demands linearly from 0 to 5 Mbps using a physical rate of 12 Mbps. Figure 6.8 shows the carried throughput measured at the application layer for all three stations. Solid lines are drawn for values obtained by simulation and the dotted lines are the values provided by our algorithm. In saturation conditions and according to the algorithm presented in Section 6.1.5.2, the throughput that any of these three stations will obtain (S_j) is about 5 Mbps, and this is the value to which each station converges, as seen in Fig. 6.8. As C increases its offered traffic, not only is it clear that less capacity is available for greedy stations, but also since C is the slowest station, the global throughput is decreased as well.

However, the carried traffic can be obtained from direct measurements at the AP.

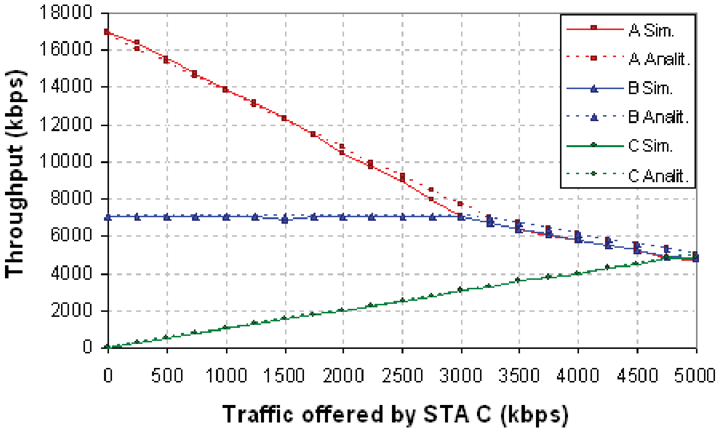

Figure 6.8: Performance anomaly captured by the AAC algorithm

The contribution of the proposed algorithm lies in its capacity to predict the potential throughput that a station will get if associated with a new AP. The next set of simulations compares the predicted AAC with the throughput carried by the newly associated station obtained through simulations. The AP is serving four saturated and mobile stations. In the beginning, the four stations are placed together, 1 m away from the AP. At time 0.0 s $STA4$ starts moving away from the AP, $STA3$ at 5.0 s, $STA2$ at 10.0 s and $STA1$ at 15.0 s (see Fig. 6.9(a)). All four STAs move at 1 m/s in opposite directions (N, S, E, and W). As they move away from the AP, the SNR is decreased and the physical rate is adapted dynamically according to the values shown in Fig. 6.7. Figure 6.9(b) shows the evolution of this rate adaptation through the simulation. The stations stop moving when they reach the border of the AP's coverage (30 m).

The algorithm is executed every 0.5 s, providing an estimation of AAC. In parallel, a new simulation is run now adding a fifth station associated with the AP at the same instants as the AAC estimation is done. Then its saturation throughput is measured twice: first, associated at 6 Mbps, and again at 48 Mbps. Estimations and simulation results are shown in Fig. 6.10. Logically, the available throughput is decreased by steps every time an associated STA adjusts its physical rate to a slower modulation, as shown in Fig. 6.9. The relative error produced by the algorithm that was observed in this scenario is always between 5% and 6%.

The next set of simulations is performed to show the accuracy of the model and the algorithm in a cell with users with varying traffic demands. We define four modes for the stations in a cell:

0. *OFF*: STA is inactive

1. *LOW*: MSDU 500 Bytes and 1 Mbps offered load

2. *MED*: MSDU 1000 Bytes and 2.5 Mbps offered load

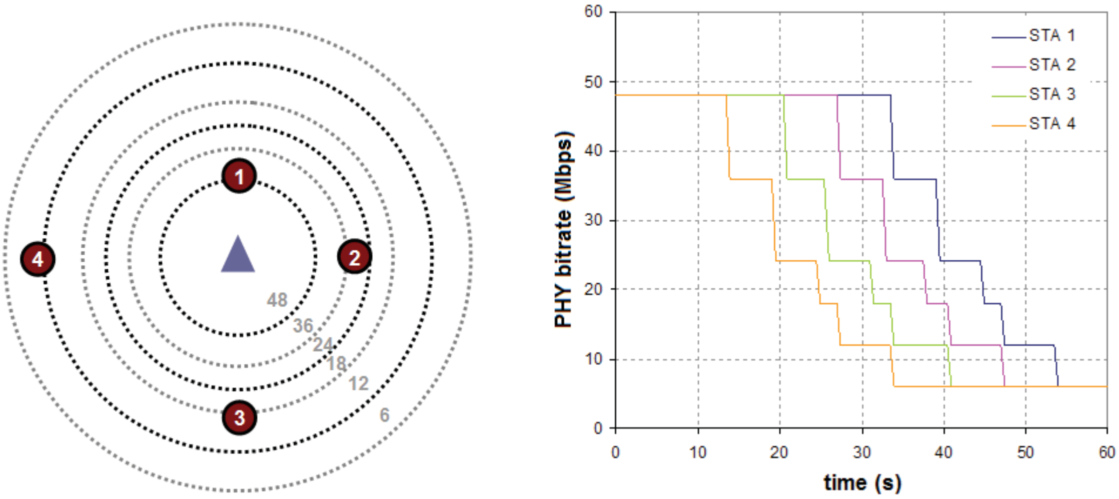

(a) Picture of the simulated scenario after 35.0 s (b) Effect of rate adaptation moving 1 m/s

Figure 6.9: Simulated scenario with mobile STAs

3. *SAT*: MSDU 1500 Bytes and saturation

In order to see the effects of different modulations, the next scenario consists of 6 stations, each of which is using a different modulation (48, 36, 24, 18, 12 and 6 Mbps); their demands vary according to the behavior shown in Fig. 6.11(a).

$STA1$ starts transmitting in mode 1 at time 0.0, $STA2$ at 0.5, etc. Figure 6.11(b) shows the AAC and simulation measurements for a seventh STA that is associated at 48 Mbps and 6 Mbps. In this case, the relative error is also in the range of 5 - 6%.

Finally, a large number of random simulations are performed with an increasing number of active users. The stations are placed randomly within the AP's coverage area and their physical rate is chosen according to Fig. 6.7. Their traffic profile is either *LOW* or *MED*. Figure 6.12 compares the average aggregate throughput obtained through simulations with the values provided by our model. This is the most complete set of simulations, since the scenarios evaluated include hidden nodes, channel errors (as a function of each transmission's SINR), different modulations (according to Fig. 6.7) and different traffic demands. Relative error is kept within small values (3 - 5%) for 20 or less stations, but increases to 9% with a very large number of users per cell.

6.1.5.3.3 Practical measurements The algorithm is fast and its implementation requires little resources. Therefore it is suitable for running on devices with limited features, such as any commercial AP. In order to prove the viability of our proposal, we implemented the algorithm in a commercial AP running a Linux-based OS: the 4G Systems AccessCube. The algorithm presented in Section 6.1.5.2 was programmed in standard C

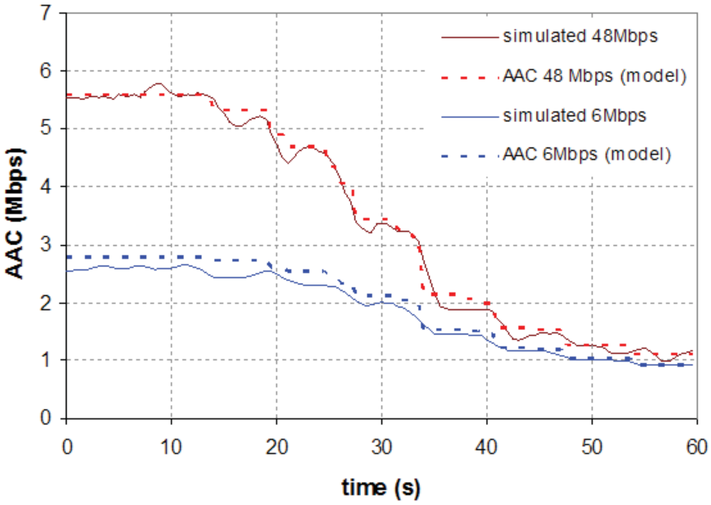

Figure 6.10: AAC estimations for a new 11g STA in a cell with four moving STAs

for Linux and cross-compiled to run on the AccessCube's MIPS architecture.
The AccessCube's OS includes the HostAP driver for the Intersil Prism 2.5 based 802.11b devices on board. This driver provides a helpful collection of statistics which are accessible to applications via the */proc* filesystem. For each associated station, the driver regularly updates information including the number of packets and the number of Bytes sent/received by that station, the number of packets sent at a given rate, and the SNR of the last received packet. The driver also provides information about the performance of the AP itself, e.g. packet discards, retransmission attempts, etc. All these parameters allow an application to fill in the variables needed to process the algorithm in a timely manner, so the effects of varying flows and channel conditions can be taken into account. The resources employed by a single run of the algorithm are insignificant in most cases. In this way, the performance of the AP is not affected by the extra load, even though the AP is serving a large amount of clients, since most of the AP functionality relies on the wireless card firmware, which does not share resources with user space applications. In the extreme that the AP has to inspect statistics of 50 associated stations, the AccessCube's CPU time usage needed to compute the capacity available for a single station ranges from 20 to 30 ms and the memory usage reported by the OS is below 1.1%

Here we present results obtained from measurements in a small testbed. The purpose of these measurements is not only to validate the algorithm in a real scenario, but also to demonstrate that the AAC estimations proposed in this section can be performed in real time by commercial APs. The testbed topology is shown in Fig. 6.13(a); there are three stations A, B and C, associated with the same IEEE 802.11b AP. All three stations have different traffic profiles: A is the source of an MPEG-2 video stream. The MPEG-2 video codec can be formatted as constant length packets of 188 Bytes (this is called Transport Stream) that can build payloads of k x 188 Bytes. The transport of TS packets over

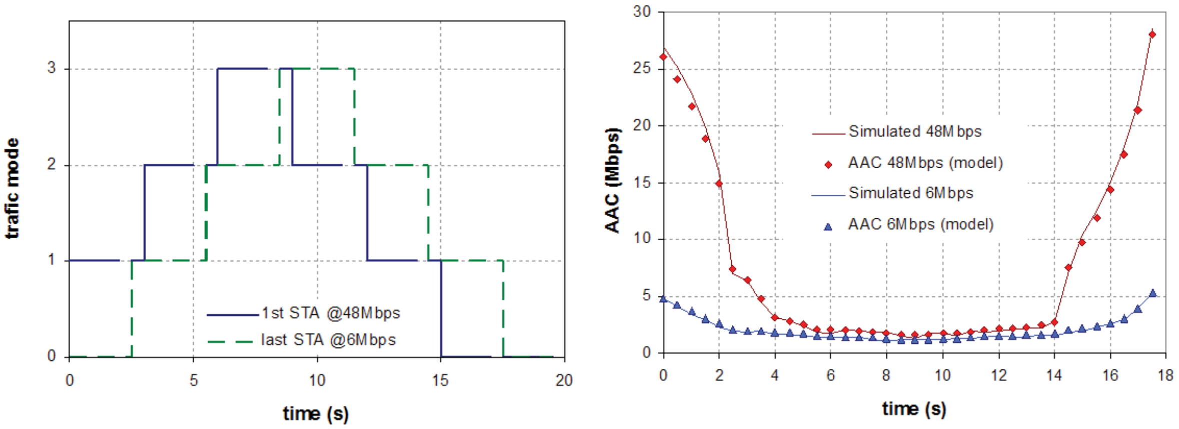

(a) Variation of a STA's traffic demands

(b) AAC estimations for a new 11g STA in a cell with 6 STAs

Figure 6.11: Simulated scenario with varying traffic demands

IP/UDP/RTP usually includes 7 TS packets = 1316+40 Bytes, in order to approach the Ethernet MTU and maximize efficiency. A's packets are spaced out so that a 2 Mbps CBR stream is obtained. In B, a greedy application is always trying to send as many 1500 Byte UDP packets as possible, while C follows a bursty pattern: the average time between consecutive bursts is 20 s, the average duration of a burst is 8 s; bursts consist of 1000 Byte UDP packets in such a way that STA C reaches saturation.

The stream originated in A is sent to $PC1$ while $PC2$ is the destination of data flows originated in B and C. $PC1$ and $PC2$ are directly connected to the AP by means of a 100 Mbps Ethernet segment, so we can guarantee that the bottleneck resides in the air interface. A sends frames at 11 Mbps, C uses 1 Mbps and B decreases its bitrate one level (modulation) every 20 s, starting with 11 Mbps. For this experiment, the AP runs the algorithm once per second to provide an estimation of the maximum capacity available for station A.

Measurements are shown in Fig. 6.13(b): during the first 10 s station A has to compete with only one element, station B, which is sending frames at 11 Mbps. The measured capacity is greater than A's requirements so A can therefore carry all its offered traffic. From 10 to 20 s, a burst of C's packets makes the capacity measurements for A fall below 2 Mbps, which corresponds to the real throughput obtained by A's flow. During the period 20-28 s, B sends frames at 5.5 Mbps, and the capacity for A measured at the AP is again greater than A's offered traffic. As can be seen in the figure, the subsequent measurements are representative of the actual throughput obtained in A's transmissions. This experiment shows that the implementation in a commercial AP of the algorithm presented in previous sections is able to detect significant capacity fluctuations affecting a given station. Note that the algorithm is slightly overestimating the actual capacity; the

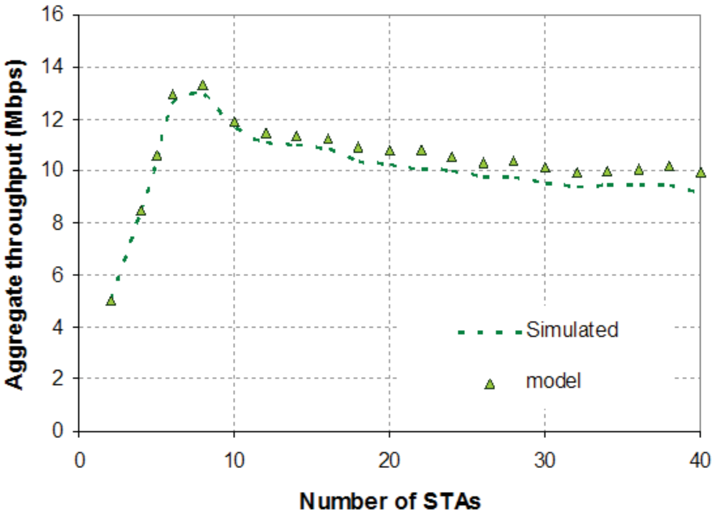

Figure 6.12: Aggregate throughput in random scenarios

values provided are, on average, less than 10% above the actual throughput measured at the same time instant.

6.1.6 Conclusions

After studying many different metrics used in IEEE 802.11 to describe the load of a cell, we conclude that none of them captured all the phenomena affecting the cell performance: the number of active users is not enough in packet switched networks since users may have different traffic profiles; packet loss can be high due to bad channel conditions, regardless of the actual load; carried traffic does not take multi-rate STAs into account; service time, MAD and busy time may not reflect the consequences of a high number of active users (when their offered traffic is low). Therefore, a new metric was needed. Our definition of the Available Admission Capacity (AAC) effectively meets all those requirements. The lower the AAC offered by an AP, the higher its load.

In this section we present a formulation set and an algorithm intended to provide capacity estimations in an IEEE 802.11 WLAN cell. The novelty of this approach is that our scheme is able to provide the AAC of a cell, that is, the load that a newly associated station could carry if associated to that AP. Moreover, we demonstrate that the AP is able to carry out these estimations by itself.

Our model is based on several assumptions that may introduce some error in the estimations. The most significant approximation appears when we consider that backoff slots are shared among active stations. While this is true under saturation conditions, the probability of two stations sharing a backoff slot decreases with decreasing load and the number of users. In order to show that the error introduced is small, we performed

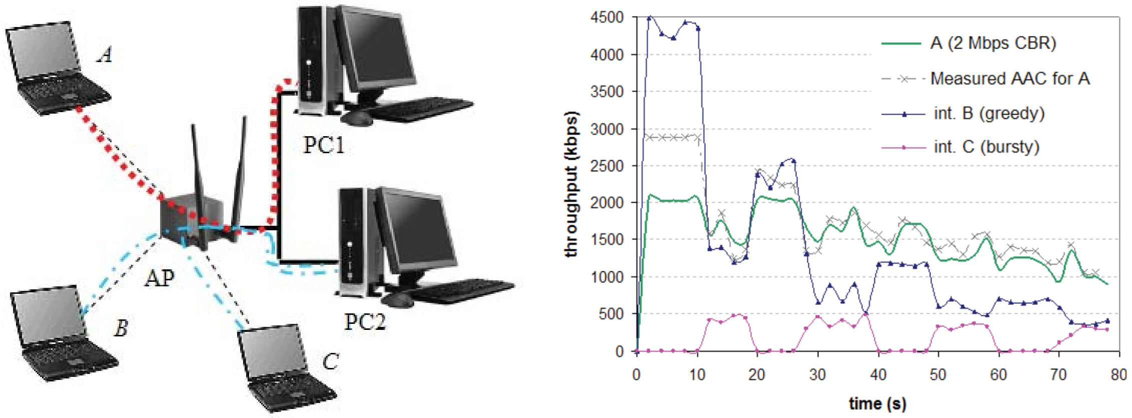

(a) Testbed for real-time AAC measurements

(b) Real-time AAC measurements

Figure 6.13: Testbed and real-time AAC measurements

a comprehensive evaluation of our approach in a wide range of scenarios. The evaluation includes analytical and simulation results and also practical measurements. Our approach is first compared with known analytical models to show that under saturation conditions (no hidden nodes) with and without channel errors, our assumptions are correct (error $<$ 1%). The simulations allowed us to test a greater variety of scenarios, including different modulations and traffic demands. The proposed algorithm produces an average relative error near 6%. Finally, the practical measurements are the proof that the algorithm can be run in real time on a commercial AP.

6.2 Load Balancing techniques

There are different techniques intended to alleviate the effects of congestion in a WLAN hot-spot like scenario. Probably, one of the most aggressive measures could be to apply highly restrictive admission control policies. For example, when the load of a given AP is reaching a pre-defined threshold level, no new users are admitted or even some of the associated users are expelled. This solution effectively reduces congestion but may produce an unfair situation, since some users with access rights are denied service. In our opinion, redistributing the load among the APs become a more attractive choice. This redistribution is achieved by applying load balancing techniques.

In this section we present two different approaches: a client-driven association scheme based on the STA's knowledge of the APs' load previously published in [128]; and a cell breathing mechanism that avoids modifying STA devices presented in [137]. Needless to say, both approaches use the same concept of load detailed in Section 6.1.5: the AAC. Recall that high loads entail low values of AAC and vice versa.

6.2.1 Related work

In addition to the load metric chosen, the load balancing techniques found in the literature can be classified according to the element in charge of taking association decisions [297]. Different approaches have been proposed in the literature that try to change the client-driven nature of IEEE 802.11 association and roaming decisions. The authors of [55], [63] and [167] propose network-controlled schemes in which client stations send the required information to a central unit, which also has access to the load information for each cell. The scheme proposed in [55] provides the best AP for association, and the network also suggests roaming to APs located further away if nearby APs are considered unable to cover the station's requirements. In order to implement these solutions, it is necessary to modify the client devices: firstly, they have to send new management frames before they are actually associated; secondly, they will no longer be responsible for association or roaming decisions. The first issue can be solved by using new radio measurements (future IEEE 802.11k devices). There is no standardized procedure for solving the second issue as yet, but it is expected to be revised by the IEEE 802.21 group, which will provide new mechanisms intended to assist handovers, and by IEEE 802.11v, which is in the development stage and will include management capabilities to allow network-directed roaming.

It is not vital to solve the second of these issues, since it is also possible to perform implicit admission control/association management. This involves actions taken on the network side that induce the desired client behavior, and therefore leave the roaming and association decisions to client stations so that hardware/software modifications are not required. In [60] the APs accept or deny new association requests depending on the respective load. When the first choice is rejected, the stations will send association requests to the next AP in the signal strength-arranged list until they are admitted. The algorithm proposed in [280] is more sophisticated but follows a similar logic. There are three possible AP states: under-loaded (will accept any request), balanced (will not accept extra load) and over-loaded (will expel the station on the assumption that it will automatically request a less loaded neighboring AP). In addition, the same authors propose in [283] to combine their load balancing with explicit admission control, based on medium access delay measurements. While these techniques provide good admission control and load balancing, they do not guarantee that all users with network access rights will be fairly served.

Another way of implementing network-directed client-driven load balancing could be achieved by using the concept of cell breathing. Cell breathing techniques consist in dynamically modifying cell dimensions by increasing or reducing transmit power. Cell breathing is a side effect in CDMA networks that reduces the cell coverage when more users are supported, but this could be advantageous in load balancing techniques if optimal strategies are applied. The concept of cell breathing for load balancing in WLANs is explained in [78]: a highly congested AP reduces its coverage radius so that the furthest stations lose connectivity and try to roam to less loaded APs. An under-utilized AP may

increase its transmit power in order to expand its coverage. Consequently, new users will roam to this AP and the load on neighboring APs will decrease. In [289], APs could even build a custom radiation pattern to balance load, but besides a very specialized RF hardware, this solution relies on the APs' perfect knowledge of their own coverage and the exact position of clients, which is hardly feasible. Reference [62] provides an in-depth analysis of cell breathing in IEEE 802.11 WLANs and proposes a centralized solution based on two different algorithms: one is aimed at reducing the load of the most congested AP and the other tries to find the min-Max load balanced solution. In [54], the authors provide five different versions of a cell breathing scheme, depending on the presence of heterogeneous or non-heterogeneous traffic demands, and depending on the availability of continuous power assignments as against a discrete set of power levels. In the case of continuous power assignment, linear programming is used to maximize throughput, while in the case of discrete values a greedy algorithm is used.
However, as stated in [136], the furthest stations may sometimes be expelled arbitrarily in spite of the fact that they may be contributing an insignificant load, depending on the applications they run. Reference [136] also gives an overview of different load balancing techniques that can be applied by using IEEE 802.11k measurements and statistics.

If the stations were able to gather more information they would be able to perform smarter associations and therefore maximize the network efficiency. In [237], clients perform different tests on all APs within range in order to determine the most suitable association. A preferable and less intrusive solution would be to introduce a trade-off between the received signal strength and cell utilization. Client devices must choose an AP that is close enough to allow frame exchange with the minimum received signal quality that guarantees correct transmission at the highest possible bit rate. However, it is advisable to avoid cell saturation by selecting the AP with the lowest load. A purely client-driven and non-intrusive association scheme is proposed in [279]: a STA observes a skewed time period of beacon frame receptions to estimate the available bandwidth. This method for estimating available bandwidth introduces large delays due to the channel observation time (several beacon intervals). In [242], the Probe Response messages sent by the APs provide the number of associated stations and information about the received signal from the requesting station. Clients use this information to calculate a weight for all reachable APs and then associate with the AP with lowest weight. The authors in [147] and [124] also propose including new load information in beacon frames so that association decisions remain on the client side. Different commercial products (e.g. [19] and [17]) use similar systems in which the APs announce their utilization in beacon frames. However, as with all proprietary solutions, there are interoperability problems that could be solved by the new IEEE 802.11k standard.

6.2.2 Client-driven approach

In the Section 6.1.5 we defined an algorithm that can be implemented in APs to provide a new load metric that is sent to requesting stations: the Available Admission Capacity (AAC), which is based on a more precise definition of an existing 11e field. An analytical study of the proposed metric was also presented in that section. In this section we describe how this metric can be applied to a client-driven load balancing approach.

This proposal was previously applied to cellular networks and makes use of the overlapping areas between neighboring cells, i.e. areas under the coverage of more than one base station (BS). Mobile stations in overlapping areas are able to reach several BSs, and when a service is required the system can decide to offer it through the BS with more available resources. The solution described in this section was first published in [128]: a new client-driven load balancing scheme based on a smart AP selection. The method uses the information available to client stations through the mechanisms and measurements provided by the new IEEE standards: 802.11e, 802.11h and, principally, 802.11k. To the best of our knowledge, this was the first practical application based on these new 802.11 functionalities, and the first time that some of the new radio measurements have been evaluated (see Section 6.1).

Based on the knowledge of the AAC-based load of the APs in range, any STA will choose an association with the AP that is able to provide the greater bandwidth to that specific STA. In this way, as explained before, load includes the effects of multi-rate STAs, carried traffic, number of users, collisions, RSSI and channel conditions. It has been shown that these measurements can be made in commercial APs at a low computational cost. The main drawback of this solution is that it requires changes on the “client side”. A third issue: how STAs acquire the new load information from the APs can be easily addressed by simply using unused fields in standard IEEE 802.11. Recall that the 11e or 11k amendments add specific informational fields intended to carry load information. The next section discusses all these implementation issues, and subsequently our proposal is evaluated against other schemes.

6.2.2.1 Implementation details

It is essential to acquire the necessary parameters when implementing the proposed AAC algorithm. The AP is the only node that is able to calculate the available bandwidth for a given user in real time, since it maintains statistics from which all of the required parameters can be derived, including the physical rate used by each station, the MSDU size, etc. Nevertheless, obtaining P_i is not trivial, since the actual number of frames sent (and lost) by its associated stations is not known by the AP without any kind of information exchange. However, the AP knows the SNR at which it receives frames from all its client nodes. From SNR values and the modulation used by the station for its transmissions, the AP is able to derive the BER (and consequently, the PER) for a certain client, which can be obtained theoretically using the formulas given in Section 3.2. The relationship between P_i and the PER is established in eq. 6.18.

The AP needs to know the amount of traffic contributed by each of its clients (TO_i) or the percentage of time devoted to each of its clients (L_i), in addition to their association rate. All of these parameters can be estimated by examining the AP's MIB counters and statistics. Please refer to Section 6.1.5 for more implementation details related to the AAC estimation algorithm.

This association management is implemented by using out-of-band signaling (through a cellular interface) which is described in the following, although we also discuss two alternatives for a more general case. The AAC calculation requires the candidate station's rate, which is not known until the association process has been completed. In the absence of out-of-band signaling, we distinguish two implementation approaches depending on whether the new station's rate is known by the AP or not, following an active or a passive scan process. In the out-of-band approach, the association control can be managed from the network side, whereas the active and passive scan solutions remain purely client-driven.

6.2.2.1.1 Active scanning As explained in Section 2.2.2, the stations can perform an active scan in order to find the APs within range. For each channel, the station willing to associate with a new AP broadcasts a *Probe Request* frame. Any AP shall respond with a *Probe Response* to the address of the station that generated the *Probe Request*. Since the AP is able to measure the SNR of the received *Probe Request*, it can predict the most suitable rate for the requesting station. The new *Probe Response* frames include the AAC field; if the AP knows the best rate that the client is likely to use in subsequent transmissions (r_i) it can calculate a specific AAC for each request. The requesting station also knows the best physical rate once it knows the signal quality perceived by the AP, which can be transmitted using the RCPI (Received Channel Power Indicator) information element present in the *Probe Response* (if 11h is enabled). The AAC is then a 2-byte codification of L_i. The requesting station i gathers all L_i values received from the APs within range and associates with the cell that will provide the best throughput ($L_i{\cdot}r_i$).

6.2.2.1.2 Passive scanning Stations that use passive scanning listen to each channel and wait for beacon frames to identify all APs within range. Unlike in active scanning, STAs do not generate any frame during a passive scan. As a result, APs are unable to calculate specific AACs for all of the possible candidate stations. Conversely, APs can set the AAC value broadcast through beacon frames with the available capacity that can be allocated to a new STA that uses the fastest possible rate. These beacons produce two important changes: STAs can select the most suitable rate by estimating the up-link margin from the TPC (Transmit Power Control) Report included in 11k beacons; and stations can determine the best AP based on the value of AAC (normalized to 1) and the number of associated stations (NSTA).

With eq. 6.19, STAs can estimate the average rate used in the cell (r_{avg}) and are able to find the specific AAC for their particular rate by applying the previous algorithm. Note that eq. 6.19 is only valid under saturation conditions, and in such conditions it provides the same values as those obtained through the active scan; otherwise, the results are inconsistent with the current available capacity. If the associated STAs do not saturate the cell or, more precisely, when the channel utilization field is approximately 1-AAC, the available capacity can be directly calculated by multiplying the broadcast AAC by the rate, although in some cases the assumptions made and the error produced affect the capacity estimation, and there is no guarantee that the best AP will be selected.

$$r_{avg} = \frac{AAC\,NSTA}{1 - AAC} r_{max} \tag{6.19}$$

This implementation provides approximate capacity estimations, so active scanning is preferred. Nevertheless, the same results could be obtained by changing the definition of beacon frames if the available capacity were calculated and broadcast for each of the rates supported by the AP.

6.2.2.1.3 Out-of-band signaling (centralized) The number of dual-mode devices (predominantly handsets) with WLAN/cellular network interfaces is growing exponentially. Future devices will be released with several network interfaces due to VoIP services and convergence driven by network operators. In this scenario, a feasible solution would be to use the cellular interface as an out-of-band signaling tool, which would enable stations to transmit information prior to WLAN association. In [244], we presented the design and implementation of a novel management architecture for a wireless mesh network. Client stations and network nodes (i.e. APs and gateways) use the cellular interface for signaling purposes only. All signaling data is transmitted to a central management entity which performs several tasks including Authentication, Authorization and Accounting (AAA), mobility management, security key distributions and dynamic frequency allocation algorithms (as detailed in Section 5.2.3).

In the case of stations, the WLAN interface is used to perform a passive scan of the medium prior to association, which gathers information on the power level, channel and BSSID of all APs within range. This information is transmitted to the management entity, which estimates the most suitable association rates that the client is likely to use with the different candidate APs, according to the power levels reported by the client. The candidate APs will compute the AAC according to the algorithm in Section 6.1.5 upon request by the central manager and send the value to the central entity. The central management entity then determines the best association for the requesting station and sends a response via the cellular interface, which includes additional details required in order to establish a successful association: the BSSID of the selected AP, channel, ESSID and encryption key. As a result, the station will be able to associate with the best AP within the client's transmission range.

The client software for association control is implemented in a Linux-based laptop (Debian 4.0 distribution), with Intel ipw3945 802.11a/b/g chipset [30] and HSDPA/UMTS PCMCIA card. The AAC module is installed on commercial APs (4G AccessCube) with an embedded Linux distribution running a 2.4.27-r11 kernel. The APs' wireless interface is an 802.11b Prism 2.5 based card, supported by HostAP driver version 0.3.7 [217], which uses the /proc filesystem to store a large number of statistics. The central entity is based on a set of software modules that run on a PC with an Intel P4 3GHz dual core processor with 2 GB of RAM and a Linux OS Debian 4.0 distribution with kernel version 2.6.18.

6.2.2.2 Evaluation

The new IEEE 802.11k APs will enable WLAN clients to select the best association based on signal strength measurements, Medium Access Delay values, Channel Busy time, number of associated STAs and available admission capacity. In our simulations we compare four different association control schemes:

- **RSSI**: current default association scheme based on the Received Signal Strength Indicator, i.e. stations select the AP that provides the strongest signal.
- **NSTA**: STAs select the AP within transmission range that serves the lowest number of associated stations. Ties between APs with an equal number of associated STAs are broken by RSSI decisions.
- **MAD**: STAs select the AP that provides the lowest MAD. Note that MAD is coded with 1 byte, the values of which will be a logarithmically-scaled representation of the current medium access delay. MAD is therefore more sensitive to low delay measurements. For example, MAD = 1 represents an access delay of between 50 and 51 μs, but if MAD = 253 the delay varies between 4396 and 5498 μs. Ties between APs broadcasting identical MAD values are broken by RSSI measurements.
- **AAC**: STAs select the AP that provides the highest AAC value following an active scan.

6.2.2.2.1 Scenario Note that current multi-cell WLANs are designed with three or four cell clusters due to spatial reuse restrictions, and only three non-overlapping channels in the 2.4 GHz band. That is to say, in a well-planned ESS most users are able to obtain services from up to four different APs. Therefore, the evaluation process we designed is based on extensive simulations in a 60x60 m square indoor scenario with four IEEE 802.11g APs, without losing generality. Thus, it is easy to scale the conclusions derived from our results to larger scenarios. We performed the evaluation by using the simulation tool

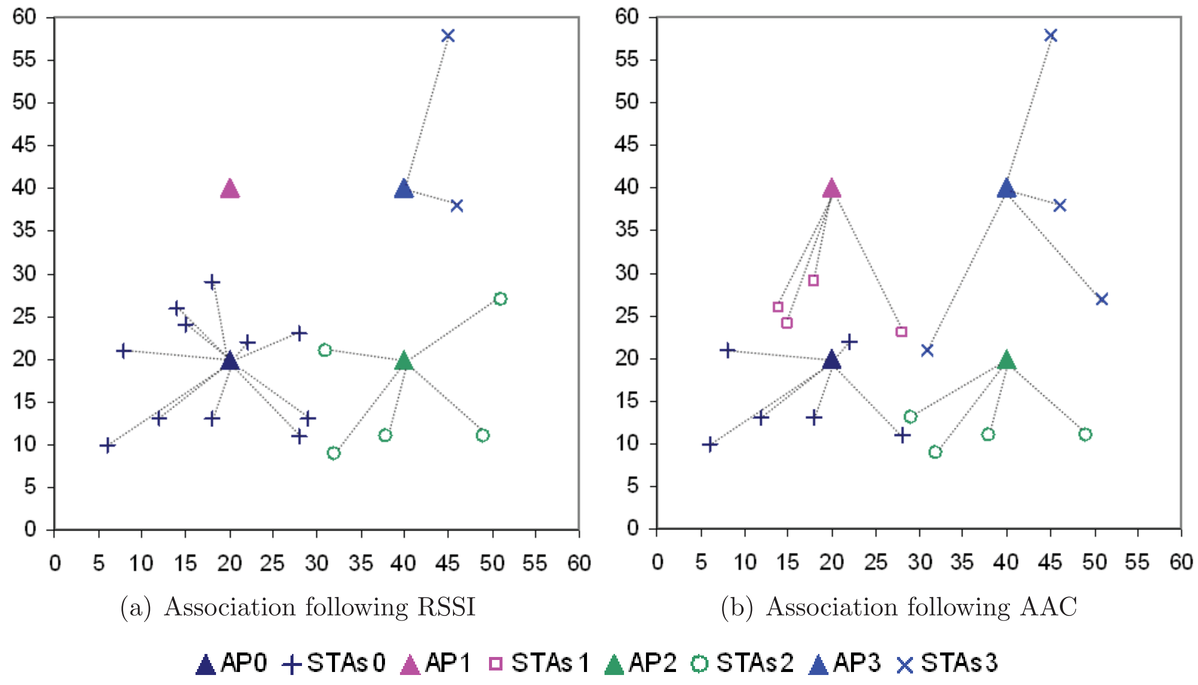

(a) Association following RSSI

(b) Association following AAC

Figure 6.14: Random distribution and association of 18 STAs

detailed in Section 3.4.2. Using this tool allowed customized simulations with a short time-to-deploy in which it was easy to define the new parameters to be measured. We ran a large number of independent simulations and obtained small confidence intervals, which are therefore not shown in the figures.

We assume that APs use non-interfering channels. As was stated in [56, 57, 158], users are static and tend to be spatially concentrated. We simulate these characteristics by placing users at random, but forcing 50% of users to be concentrated in an area of 30x30 m around one of the APs ($AP0$). This ensures that it is possible to identify the benefits of association management beyond signal strength measurements, but also a realistic scenario is met. Figure 6.14(a) shows the possible user distribution and association when the RSSI scheme is used; given the same user distribution across the scenario, Fig. 6.14(b) shows the association according to AAC.

The physical rate used for transmissions depends on the distance between a STA and its selected AP. Figure 6.7 (pag. 132) shows the correlation between rate and distance when using the propagation model for semi-open offices, detailed in Section 3.1.

6.2.2.2.2 Saturation The first set of simulations are intended to recreate the worst-case scenario, in which all stations are running greedy applications and there is always a frame ready to be sent in every station's transmission buffer. If a fixed rate is set for all the

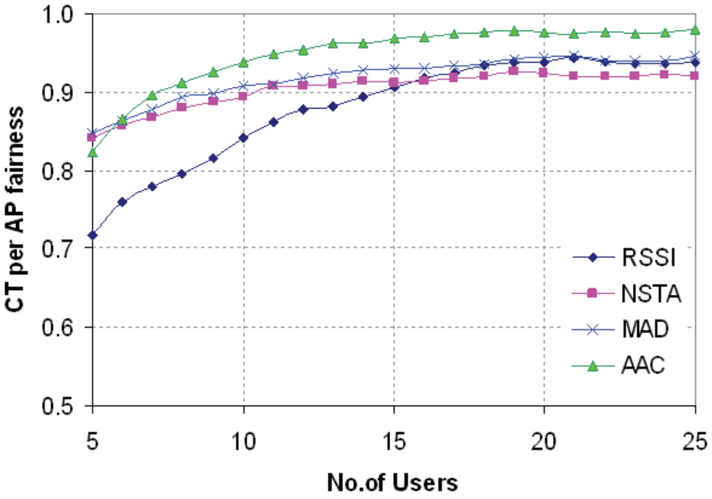

Figure 6.15: Throughput per AP fairness under saturation conditions

stations, there will be unnoticeable differences between the throughput carried by the APs, but if we use MAD, Service Time or AAC to measure the load, an unbalanced situation is revealed: 50% of the users are associated with $AP0$ when the traditional RSSI approach is used, which is clearly the worst solution. In this specific case (saturation and fixed rate), NSTA associations are slightly better than those using MAD and AAC, since all stations represent the same load; i.e. balancing the number of stations implies balancing load. The effect of load balancing can be seen when multiple rates are added, even if we measure carried throughput (CT). Figure 6.15 shows the throughput balancing fairness between the APs. Our approach improves the traditional RSSI scheme by 15%. Fairness is measured by using the known Jain's Index [168]: β is a value between 0 (unfair) and 1 (fair); if only k of n flows receive equal resources and others receive none, the index is k/n. S_i is the throughput achieved by STA i.

$$\beta = \frac{\left(\sum_i^n S_i\right)^2}{n \ \sum_i^n S_i^2}; \quad 0 \leq \beta \leq 1 \tag{6.20}$$

However, it is easier to identify the benefits of our approach if we analyze the load from the user perspective. Figure 6.16(a) shows the aggregate throughput of the whole network as we increase the number of associated stations. It can be seen that the RSSI-based scheme has the best performance, since it is the only approach that can guarantee that all clients will always use the highest possible physical rate. In contrast, RSSI associations give an unfair share of resources among users. Figure 6.16(b) shows that the RSSI-based association scheme provides an imbalanced distribution of available throughput, and is by far the most unfair solution (AAC scheme improves RSSI by more than 45% in the worst case); conversely, MAD provides a high degree of fairness but low throughput.

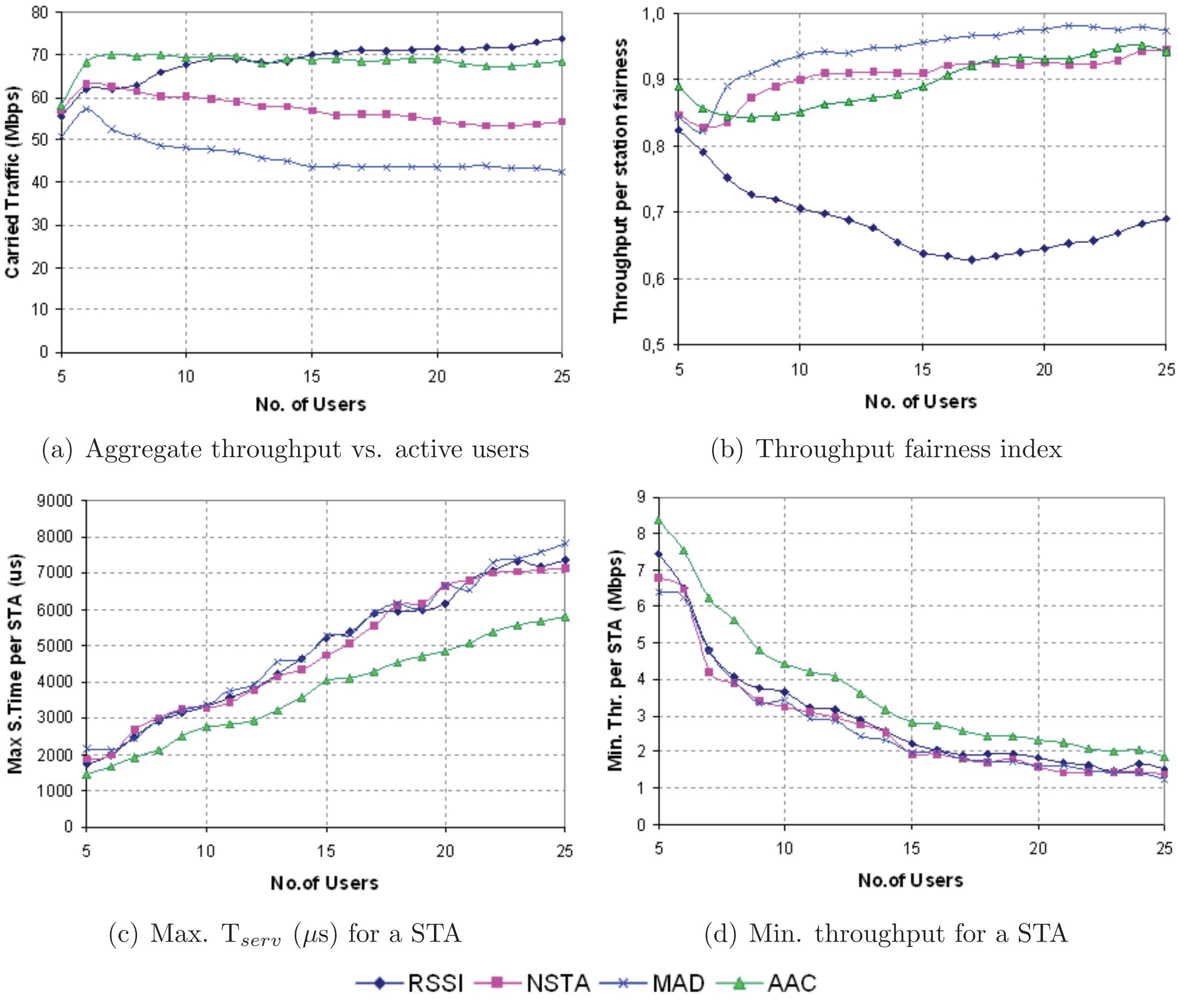

(a) Aggregate throughput vs. active users

(b) Throughput fairness index

(c) Max. T_{serv} (μs) for a STA

(d) Min. throughput for a STA

Figure 6.16: Evaluation of different association schemes under saturation conditions

Two useful parameters for measuring performance and fairness are the maximum service time and the minimum throughput obtained by a given station (see Fig. 6.16(c) and 6.16(d) respectively). By monitoring the evolution of these parameters, it can be seen that our AAC proposal provides a clear improvement over the other solutions (up to 25% and 35% respectively over RSSI). After reviewing this set of figures, we can conclude that AAC has both a high fairness index and a good aggregate throughput under saturation conditions.

6.2.2.2.3 Degree of saturation The next set of simulations was run under the same conditions as described in the previous subsection, with the exception of user traffic load profiles. In this case we define three different user types according to level 2 traffic

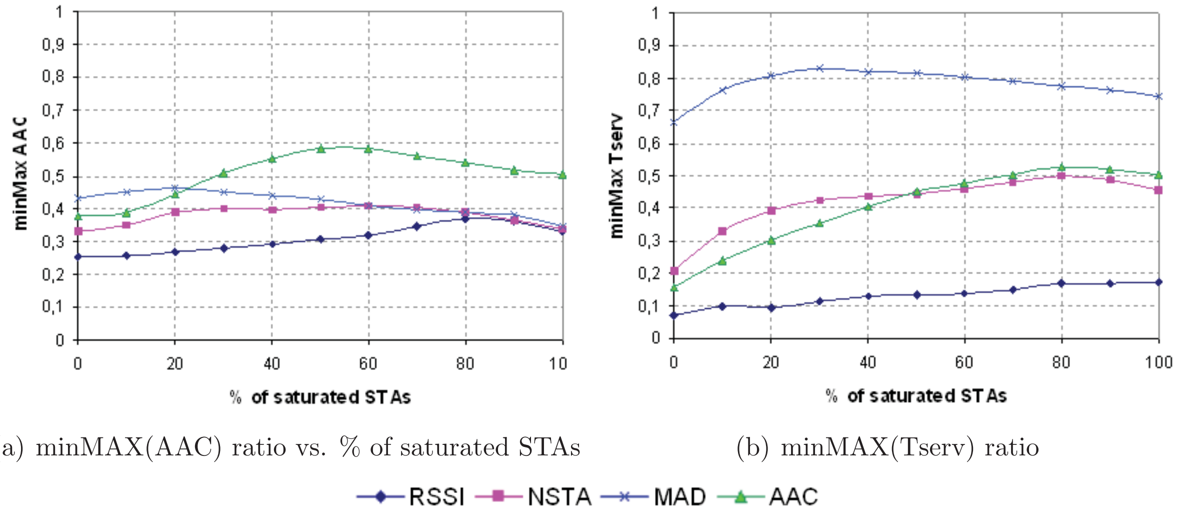

(a) minMAX(AAC) ratio vs. % of saturated STAs

(b) minMAX(Tserv) ratio

Figure 6.17: AAC and T$_{serv}$ optimization through association control

demands:

- **Saturated**: greedy users. They always have frames ready to be sent (frame size: 1450 bytes).

- **Medium**: bursty traffic of up to 2.5 Mbps with an average frame size of 1024 Bytes. This user type accounts for 70% of non-saturated users.

- **Low**: constant bit rate of up to 1 Mbps with a frame size of 500 Bytes. This user type accounts for 30% of non-saturated users.

Since the differences in carried throughput between the four schemes are minimal, it is interesting to show the level of balancing achieved with AAC (Fig. 6.17(a)) and MAD (Fig. 6.17(b)). Both graphs show the minMAX ratio (the lowest value found on an AP divided by the highest). It can be seen that RSSI provides the worst level of balancing in all cases. The load (understood as the proportion of time during which the cell is busy) which is translated into available capacity is better balanced when AAC is used, but it is outperformed by MAD if we measure T_{serv}. However, the best balance measured in terms of T_{serv} does not compensate for the poor performance of MAD if we also consider the network capacity. Figures 6.18(a) and 6.18(b) show the aggregate throughput and the fairness index as the number of saturated stations is increased (from 0 to 100%). The Fairness Index has been redefined to include non-saturated stations. Note that the presence of *Low* and *Medium* stations with eq. 6.20 would produce low fairness values even though the stations are able to carry all of their offered traffic. The new fairness

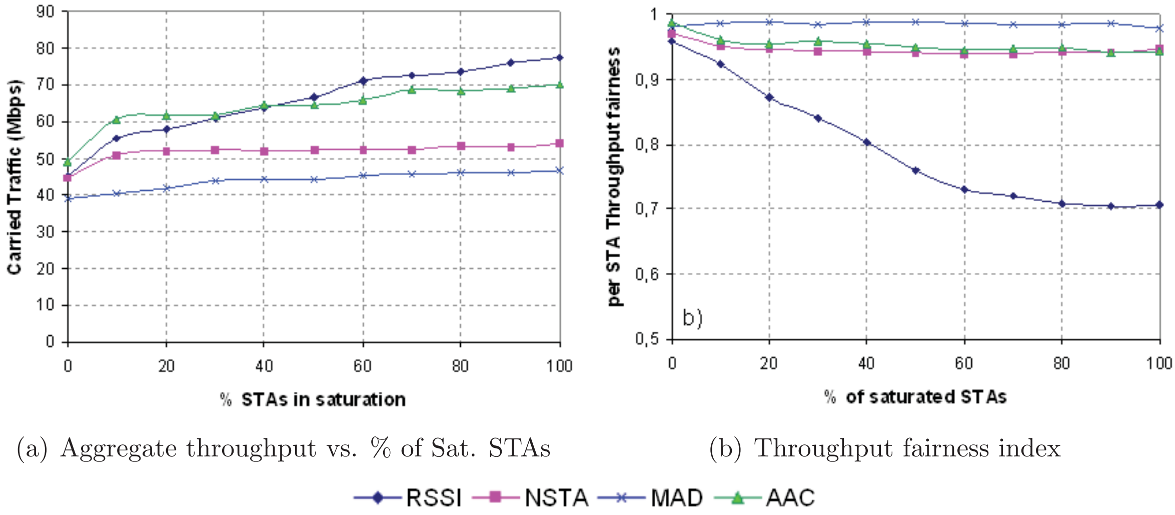

(a) Aggregate throughput vs. % of Sat. STAs

(b) Throughput fairness index

Figure 6.18: Throughput optimization through association control

index is calculated as follows:

$$\beta = \gamma \frac{\left(\sum_i^{i \in Sat} S_i\right)^2}{n_s \sum_i^{i \in Sat} S_i^2} + (1-\gamma) \frac{\left(\sum_i^{i \notin Sat} S_i/OT_i\right)^2}{(n-n_s) \sum_i^{i \notin Sat} (S_i/OT_i)^2} \tag{6.21}$$

where γ is the proportion of saturated stations, S_i is the throughput obtained by station i and OT_i is the traffic offered by station i; n_s is the number of saturated STAs.

Although the differences between the various association schemes become clearer as the number of saturated stations increases, we can draw almost the same conclusions as under saturation conditions: RSSI scheme provides the best aggregate throughput at the cost of providing the poorest fairness. AAC still provides good aggregate throughput (the best for a small number of saturated STAs), while at the same time fairness is assured.

We argue that if the conclusions are the same as those for saturation conditions, there is no need to apply more complex traffic generators. The use of more realistic traffic patterns mainly affects the way in which new radio measurements are taken: the key parameters that must be taken into account are the duration of the measurement, the frequency with which they are taken, the number of repetitions and the time during which the results remain valid, depending on the precision required. The authors of [218] suggest using confidence intervals to optimize these parameters. See also reference [240].

6.2.2.2.4 Fairness vs. Throughput The AAC proposal could be improved in terms of fairness. We propose a weighted combination of MAD and AAC metrics in order to

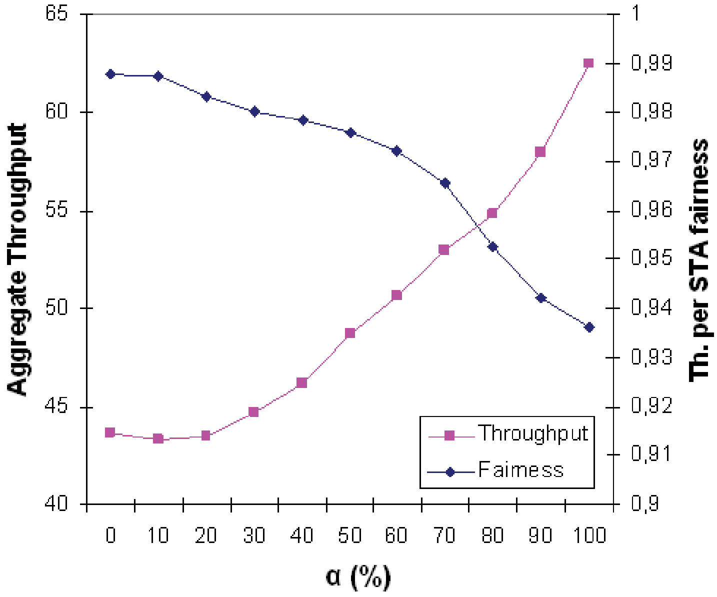

Figure 6.19: Throughput vs. Fairness with different values of α

resolve the throughput/fairness trade-off: a STA compares AAC and MAD values received from all APs in range; if $\text{AAC}_j > \alpha\text{AAC}_{j-1}$ and $\text{MAD}_j < \text{MAD}_{j-1}$, then the STA selects AP_j for association. Where $\alpha \in [0, 1]$, a value of 0 produces a pure MAD selection and a value of 1 is based only on AAC. Figure 6.19 shows that the aggregate throughput decreases as the fairness index is increased in the same scenario described in 6.2.2.2.3, with 25 users and 40% of saturated STAs. Fairness decreases rapidly when $\alpha > 0.7$, but a slight loss of fairness in this region leads to a dramatic improvement in throughput: for example, a 3% decrease in fairness is compensated by a 10% improvement in throughput.

6.2.2.2.5 Testbed Finally, we present practical measurements taken in a small testbed that can be used as a proof of concept. These practical results agree with the conclusions derived from the more comprehensive simulation study previously detailed. The scenario consists of two IEEE 802.11b APs, $A1$ and $A2$. $A1$ serves one STA at 1 Mbps; $A2$ serves two STAs, both using 11 Mbps. All STAs are in saturation, sending 1500 Byte UDP frames. A fourth STA is activated, which receives signal from $A1$ with SNR = 20 dB and from $A2$ with SNR = 11 dB; the new STA can be associated with $A1$ at 11 Mbps, while it should use 5.5 Mbps if associated with $A2$. It is clear that following an association based on RSSI, the number of associated stations, or either based on carried traffic, the new station chooses $A1$. In saturation, competing against one STA at 1 Mbps, the new client gets 766±30 kbps. However, if the new station associates with $A2$, the throughput is increased to 1672±54 kbps, even though it has halved its physical rate. Those numbers correspond to the values of AAC: 766 and 1672 is the available throughput for a new station in $A1$ and $A2$ respectively. Therefore, according to our proposed scheme, the new STA chooses $A2$ obtaining more resources. Only when the three previously associated STAs offer less than 400 kbps each (i.e. not in saturation), all four schemes coincide in selecting $A1$.

6.2.3 Cell Breathing

In the previous subsections we have described the issues derived from the default STA $\leftrightarrow$ AP association scheme used in IEEE 802.11 WLANs: although a Hot Spot could be served by several APs, most users will be connected through the AP that provides the strongest signal. As an inherent consequence, over-loaded APs offer users in congested areas a very low QoS, while nearby APs remain under-utilized. This behavior is determined by the roaming process. Thus, at this stage it is necessary to recall the previously mentioned mechanism (detailed in Section 2.2.2.2).

A STA keeps track of the Beacon frames received from its current AP. When the quality of beacons drops below the cell search threshold ($10 < CS_{Th} < 30$ dB), the STA initiates an active scan and sends out Probe Request messages on all the available channels. The APs receiving the Request will send a Probe Response back. When an AP is found whose responses improve the current AP's Beacons quality by at least ΔSNR (usually $6 \leq \Delta SNR \leq 8$ dB), the STA initiates a cell switch. If a better candidate is not found, the STA returns to the current AP's channel and the scan sweep is repeated periodically.

There are different solutions to the unfair situation explained above, the most evident of which consists of using a different roaming criterion (e.g. AP load), as proposed in the previous section, but it requires some changes in client devices. We now choose to keep any modification transparent to the end user, who can be equipped with off-the-shelf devices. For these reasons, in [137] we proposed a new AP-driven load balancing scheme based on cell breathing, intended to alleviate the congestion in hot spots: congested APs reduce the size of their cells; alternatively, under-utilized APs increase their cells to attract further stations. In our approach, neighboring APs cooperate in order to improve performance and fairness levels. With this aim, APs can make use of the information available to client stations through the mechanisms provided by the new IEEE standards: 802.11e, 802.11h and, principally, 802.11k. We can consider this solution as "network-directed", but it still respects the client-driven nature of standard IEEE 802.11 association policy.

In this section we describe a distributed algorithm, based on cell breathing, with which the APs in an IEEE 802.11 ESS WLAN are able to tune their cell size according to their load and also to their neighbor's load. As shown in the performance evaluation, this technique improves the fairness and the performance levels from the APs' point of view, but also from the end users' perspective [137].

The concept of load is inherited from the previous sections. We consider that the capacity that an AP is able to offer to a new STA (AAC) is the best load metric, since it captures all phenomena affecting the performance of an AP. Refer to Section 6.1.5 for more details on the definition and estimation of the AAC.

6.2.3.1 Distributed load balancing algorithm with cell breathing

First of all, we have to distinguish between two independent types of transmitted power management: cell dimension management (Cell Breathing) and transmitted power control (TPC). As previously explained, Cell Breathing tries to improve load balancing among neighboring APs, while TPC is aimed at reducing power consumption, interference and the near-far effect [251].

From the client station's point of view, the cell dimensions are determined by the energy of received *Beacon* frames and *Probe Responses*. Then, an AP can set its optimal cell dimension so that the farthest client that the AP must serve receives Beacons with SNR $> CS_{Th}$. However, as in [62], the power used to transmit data frames can be higher so that the user's experience is not degraded. Hence, an optimum TPC algorithm is assumed (e.g. see [253, 255]) for the exchange of data frames between an AP and its clients, using the minimum power that does not degrade the performance of the communication. For this reason we distinguish between tx range (determined by the maximum transmission power allowed) and cell size (determined by beacons and Probe Responses).

In our approach, APs are responsible for computing their own load and let their neighbors know about it by either periodic or triggered updates. Similar to [280], APs can be in one of the following three states, according to their load, as compared with their neighbors':

- **Fair**: the AP's load is similar to the average load in the neighborhood. An AP in this state will not take any action regardless of its neighbor's behavior.

- **Gull**: the AP's load is larger than the average load in the neighborhood. An AP in this state is willing to reduce its cell and will try ask its neighbors for help.

- **Willing**: the AP's load is below the average load in the neighborhood. An AP in this state is willing to increase its cell in response to a neighbor's appeal.

In order to determine the AP's load we propose the AAC metric, defined as the capacity available for a new station that uses the fastest modulation (see Section 6.1.5. Logically, as the congestion increases, the APs' AAC decreases. Then, we consider that an AP i is *Gull* if $AAC_i < \overline{AAC_i} - \delta$, where AAC_i is the capacity available in AP i, $\overline{AAC_i}$ is the average capacity in i's neighborhood and δ a threshold value used to add hysteresis, thus improving the stability of the system. Analogously, an AP i is *Willing* if $AAC_i > \overline{AAC_i} + \delta$. If none of the previous conditions is met, the AP is in state *Fair*.

The value of δ is set dynamically according to $\overline{AAC_i}/3$ (i.e. 33% of the average load). Fairness worsens as δ increases, but a small value of δ produces an unstable system. The optimum value was chosen after a simulation-based study. This 33% is the lowest value that guarantees convergence. Finally, two APs are considered neighbors if there is at least one client within transmission range of both APs.

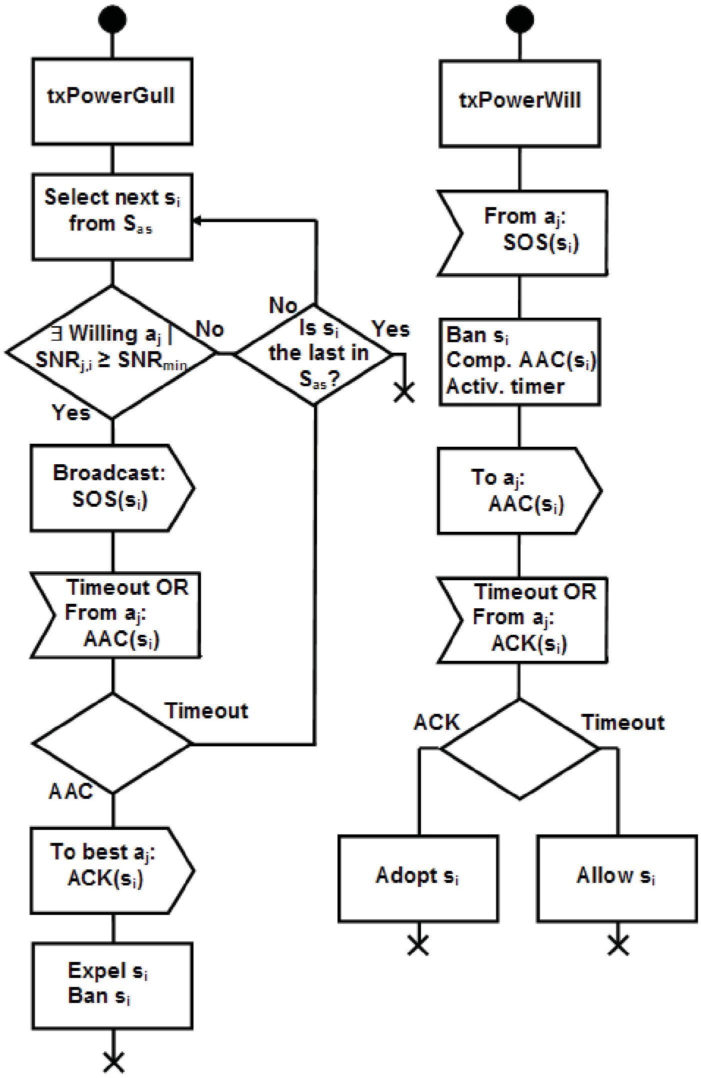
txPowerGull
Select next s_i from S_{as}
$\exists$ Willing a_j | $SNR_{j,i} \geq SNR_{min}$
No
Yes
No
Is s_i the last in S_{as}?
Yes
Broadcast: SOS(s_i)
Timeout OR From a_j: AAC(s_i)
Timeout
AAC
To best a_j: ACK(s_i)
Expel s_i Ban s_i
txPowerWill
From a_j: SOS(s_i)
Ban s_i Comp. AAC(s_i) Activ. timer
To a_j: AAC(s_i)
Timeout OR From a_j: ACK(s_i)
ACK
Timeout
Adopt s_i
Allow s_i

Figure 6.20: Behavior of a) a *Gull* AP, and b) a *Willing* AP

The behavior of *Gull* and *Willing* APs is detailed in Figure 6.20. Note that initial and final states are connected and that the process can be interrupted if the state of the AP changes. The first action taken by an AP is to arrange its cell size according to its state. *Fair* and *Gull* APs will run *txPowerGull()* to reduce the size of their cells so that the client receiving the poorest signal detects beacons with SNR > CS_{Th}, or the minimum transmission power is reached (see alg. 4 on pag. 160); *Fair* APs will not take any further action. A *Willing* AP will run *txPowerWill()* increasing its cell size so that none of the STAs associated to a neighboring AP roams to it, or the maximum transmission power is reached (see alg. 3 in pag. 160).

We define S as the set of STAs (s_i) and A the set of APs (a_j); S_{as} is the subset of the elements of S that contains the STAs associated to a given AP a_j, and S_{rg} is the list of STAs within a_j's range. Both lists are arranged in decreasing order according to the SNR computed from a_j's beacons. $SNR_{j,i}$ is the SNR of a_j's beacons as seen from STA s_i, while SNR_i is the SNR of the beacons that s_i receives from its current AP. Then, for any a_j:

$$S_{rg}^{j} = \{s_i \forall i \in S \parallel SNR_{j,i} > SNR_{min}\} \tag{6.22}$$

$$S_{as}^{j} = \{s_i \forall i \in S \parallel SNR_{j,i} > SNR_{k,i} \forall k \in A\} \tag{6.23}$$

We assume that S_{as} and S_{rg} are always updated thanks to the complete collection of statistics provided by an independent process (see next section). A *Gull* AP will then select the first s_i from its S_{as} that is able to roam to a *Willing* AP. The *Gull* AP will send an *SOS* message to all APs within range of s_i. A *Willing* AP receiving a *SOS* message will compute the AAC value for that particular STA and will forward this value to the requesting *Gull* AP. This AAC is computed by taking into account the fact that an optimal TPC is used for data exchange. The *Gull* AP then sends an acknowledgment only to the best AP candidate and adjusts its cell size expelling the selected STA (s_i). In turn, the adoptive AP adjusts its cell size to accommodate s_i. In order to avoid undesired handovers, all APs receiving a *SOS* message will ban the announced s_i (e.g. via a black list of MAC addresses) until the process is complete.

6.2.3.2 Implementation issues

Although the algorithm detailed in the previous section has not been fully implemented as yet, some of the functionalities required have been previously tested in real scenarios. For example, the signaling required to communicate neighboring APs could be easily carried out by means of a common wired backbone. If there is no such common backbone, APs could still participate in the distributed algorithm using a wireless distribution system based on mesh concepts as proposed in [127] and discussed in Section 7.3.2.

As in [62] and [253], one of the requirements for the APs is the possibility to set the transmission power in a per-packet manner. In this way, APs can arrange their cell size adjusting the transmitted power for Beacons and Probe Responses and at the same time running an effective TPC for data. The AAC computation adds another requirement for the APs: an updated collection of statistics is required at application level in order to allow the AAC estimation in real time. The implementation issues regarding AAC computation are detailed in Section 6.1.5.

Nevertheless, the information needed by the APs to run the algorithm described in the previous section represents the main implementation issue. Any AP should know the complete list of STAs within transmission range and the list of APs that any of these STAs can reach, including SNR of beacons and potential SNR for data. The latest IEEE 802.11 standard [24], which includes the 802.11h amendments, together with the new 802.11k standard will ease the acquisition of this information, as detailed next.

The potential SNR for data exchange between an AP and all the client STAs in range can be obtained by means of an 11h's TPC Request/Report or an 11k's Link Measurement Request/Report. These two mechanisms are similar and allow the estimation of the link margins between two stations. The requesting AP announces the transmitted power used to send the request (maximum allowed transmitted power) and the requested STA responds with the link margin according to the SNR of the received request. The response also includes the transmitted power used to send the frame. APs are also able to retrieve information about the SNR of received beacons using the Beacon Request/Report defined in 802.11k. A STA receiving a Beacon Request will respond with a Report containing statistics, including SNR, power, channel and BSSID, of received Beacons and Probe Responses.

The AP still has to know the potential SNR for data between its in-range STAs and the neighboring APs. This could be solved either adding an extra signaling among APs or independently, using the 802.11k frames: Measurement Pilot. Similarly to Beacons, these frames are transmitted pseudo-periodically by APs at a small interval, but a Measurement Pilot is smaller than a Beacon and is transmitted more often than a Beacon. STAs also include statistics of received Measurement Pilots in Beacon Reports, so, if APs send these frames at the maximum allowed transmission power, APs could finally gather all required information.

However, our approach is also feasible with no 802.11k enabled devices. We have to note that in this case, many of the parameters can only be approximated and that STAs are required to perform active scans. In this way, all APs within the STA's range are able to obtain the uplink margin from Probe Request messages, and thus estimate the downlink margin assuming that the path is symmetric and that the power used to send the Probe message is known (max. allowed power). These assumptions also allow APs to estimate the power of received Beacons (knowing the power of transmitted Beacons and the estimated path loss). Furthermore, more signaling is required to exchange this information among APs.

6.2.3.3 Performance evaluation

This section describes the evaluation process followed to obtain the network performance in a hot-spot like scenario. The evaluation will show the improvements achieved when the load balancing mechanism is enabled.

6.2.3.3.1 Scenario The evaluation process we designed is based on extensive simulations in a 380x380 m square indoor scenario with 16 IEEE 802.11b APs evenly distributed. We ran a large number of independent simulations and obtained small confidence intervals, which are therefore not shown in the figures. The throughput carried by an AP and the throughput available to the STAs is computed by using the AAC algorithm (c.f. Section 6.1.5).

Using a path loss $PL(d) = 40 - 33log(d)$, where d is the distance between a transmitter and a receiver, $ptxMax = 15$ dBm and $ptxMin = 10$ dBm (highest and lowest allowed transmission power), we assume that with all APs transmitting at $ptxMin$, there is no coverage gap in the scenario, and that transmitting at $ptxMax$ no co-channel interference is produced (using a 4-coloring scheme). As stated in [158], users are static and tend to be spatially concentrated. We simulate these characteristics by placing users at random, but forcing a given percentage of users, $c\%$, to be concentrated in a randomly selected area of 100x100 m. This ensures that a realistic scenario is met. The physical rate used for data transmissions depends on the distance between a STA and its selected AP: if $d < 46$ m, rate = 11 Mbps; if $d < 61$ m, rate = 5.5 Mbps; if $d < 75$ m, rate= 2 Mbps; and if $d < 92$ m, rate = 1 Mbps. For $d = 92$ m $SNR_{min} = 1$ dB is not met. Finally $CS_{Th} = 20$ dB and ΔSNR = 7 dB.

Our approach (Distributed Cell Breathing - *DCB*) is compared against different mechanisms. The Centralized approach (*CCB*) of [62] is used as a reference since, as we understand, a complete knowledge of the scenario will allow better assignments. Both DCB and CCB use AAC as the load metric. We call *TPC* the solution that implements solely an optimal TPC for data exchange, but that keeps the size of the cells fixed. Finally, the *default* behavior of current IEEE WLANs is also represented in the simulations.

6.2.3.3.2 Simulation results The first conclusion derived from the simulations is that the proposed algorithm converges rapidly in the scenario depicted in the previous subsection. Then we measure the aggregate throughput in different situations. Figure 6.21(a) and 6.21(b) are obtained in saturation conditions, that is, all STAs always have buffered frames (1500 Bytes) ready for transmission. Figure 6.21(a) shows the effects of increasing the concentration (c %) with a fixed number of users (65), while 6.21(b) has a fixed c (55%) and a varying number of users. It is not surprising that the TPC solution presents the best results, since it always guarantees that all STAs use the best possible rate. In the case where STAs have different traffic profiles (packet size from 500 to 1000 Bytes and offered load ranging from 0.2 to 2 Mbps), DCB outperforms the other approaches (see Figure 6.21(c), but as the number of saturated users increase, it becomes slightly worse. The aggregate throughput has a maximum when there are 3 or 4 STAs per AP, and decreases with more users due to the increasing collision probability. Logically, as c increases (more users use less APs), the aggregate throughput decreases.

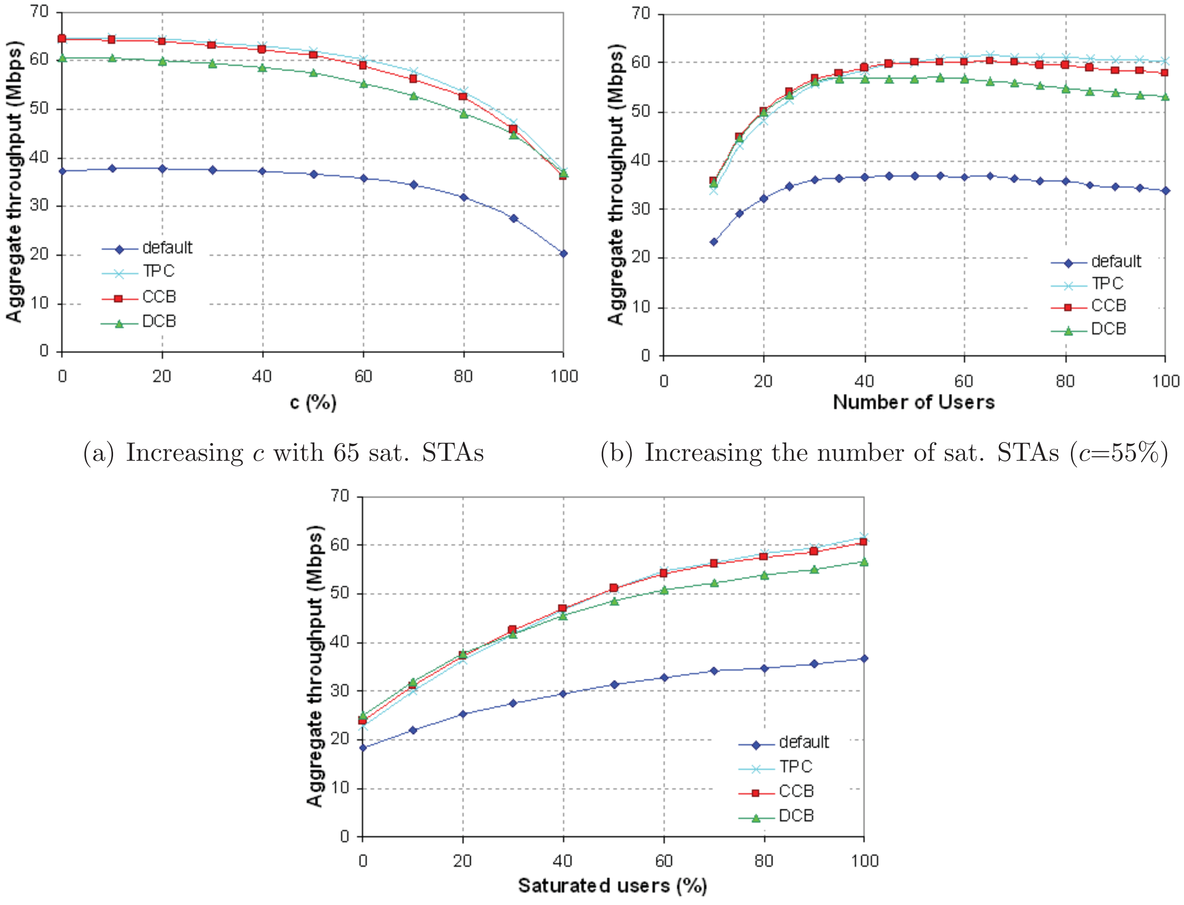

(a) Increasing c with 65 sat. STAs

(b) Increasing the number of sat. STAs (c=55%)

(c) Increasing offered traffic (c=55%, 65 STAs)

Figure 6.21: Aggregated throughput in the presence of load balancing

However, a maximized aggregate throughput does not involve that the throughput of all STAs is maximized. For this reason we also measure the fairness degree among STAs and among APs. Fairness is measured by using the known Jain's Index: β is a value between 0 (unfair) and 1 (fair); if only k of n flows receive equal resources and others receive none, the index is k/n (see eq. 6.20).

When we measure fairness among STAs, S_i represents the traffic carried by STA i and n is the number of STAs. When we measure fairness among APs, S_i is substituted by AP i's AAC, and n is the number of APs. We observe that DCB presents the best fairness values in all the cases, regardless of the number of users, the value of c, or number of STAs in saturation (e.g. see Figure 6.22). Another measure of fairness can be provided by measuring the smallest AAC announced by an AP, and the lower throughput carried by a STA. In this case, since CCB is designed to maximize AAC_{min}, its results are logically the best (see Fig. 6.23(a)). But although CCB also provides the highest minimum carried

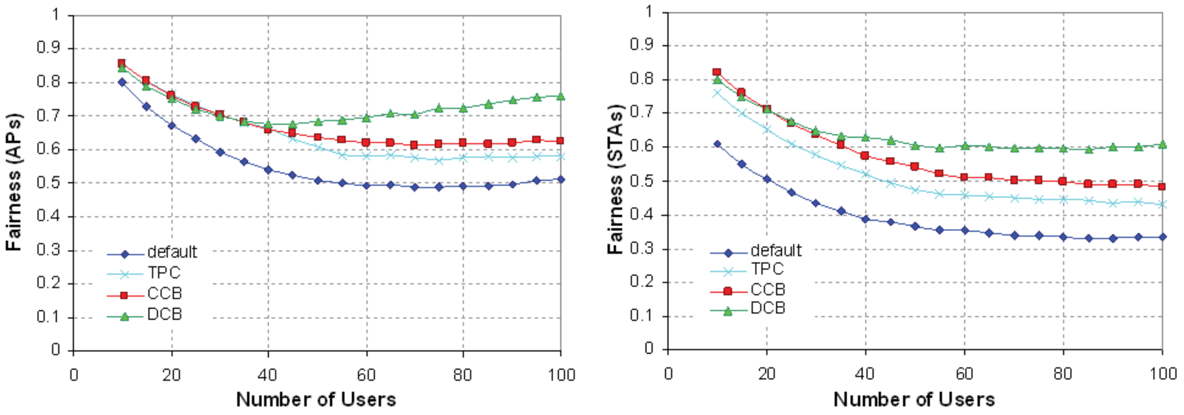

(a) Jain's fairness index among APs

(b) Jain's fairness index among STAs

Figure 6.22: Fairness among APs and STAs with increasing number of saturated users

throughput on average, (as shown in Fig. 6.23(b)), we have to note that DCB provides the best results in most simulations.

6.2.4 Conclusions

The radio measurements and mechanisms introduced by the new IEEE 802.11 standards provide valuable information that can be used by stations when they have to select the most suitable AP for association. We introduce a new load metric that can be easily estimated by the APs themselves in real-time. The new metric proposed is derived from a more precise definition of the Available Admission Capacity field and takes into account the influence of the main factors that affect the load of a WLAN: the number of users in a cell, the user rate (modulation), the signal quality and the offered traffic.

In the worst-case scenario, which is in fact a realistic situation, a large number of users are concentrated in a small area and traditional association schemes based on RSSI measurements guarantee that the fastest physical rate is reached. However, this leads to an uneven distribution of load, since most users will be associated with a small number of APs and will therefore receive a poor service. This problem can be alleviated by means of different load balancing techniques. After an in-depth state-of-the-art review, we propose two different approaches: a client-driven association scheme based on the STA's knowledge of the APs' load; and a cell breathing mechanism that avoids modifying STA devices.

In the first approach, the APs broadcast their load (AAC) in Beacon and Probe Response frames. In this way, stations are still able to perform client-driven associations and roaming in accordance with the operations defined in the IEEE standards. The new AP selection method not only provides good network performance but also ensures an even share of bandwidth among clients and a balanced load among APs. In other words,

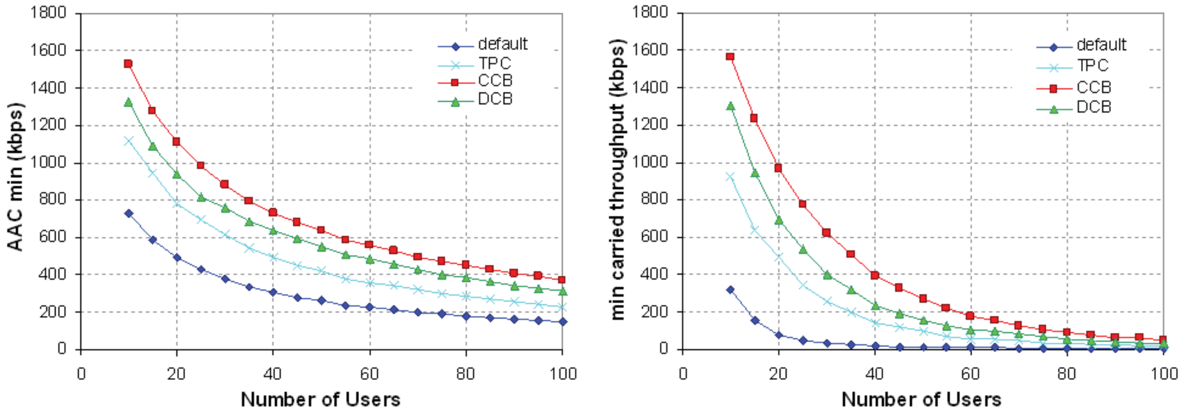

(a) minimum load (AAC) of an AP

(b) minimum throughput carried by a STA

Figure 6.23: Other metrics for showing fairness

an association scheme based on available capacity ensures load balancing among the APs of a given WLAN, increases the overall capacity and maintains fair resource sharing from the user perspective. Our mechanisms are of great help in the provision of QoS, but additional features are still required to guarantee any quality.

Moreover, in order to avoid any changes in client devices, we propose a different approach to improve load balancing using a network-directed mechanism. We present a new distributed load balancing algorithm for IEEE 802.11 WLANs, based on the idea of cell breathing. In this second approach, the APs have the ability to cooperate in order to redistribute the load among neighboring cells, in a way that is transparent to the end user, who can be equipped with standard devices. The most obvious conclusion that can be derived from the evaluation is that the absence of any kind of power control reduces the potential capacity of the network drastically. Applying an optimal TPC for data exchange ensures a better utilization of the resources, and therefore the performance of the network is improved. However, in scenarios with a high density of nodes, the average user experience can be further improved, and the congestion level on APs alleviated, if we introduce the ability to dynamically change the cell size according to the environment. Our approach not only provides good network performance but also ensures an even share of bandwidth among clients and a balanced load among APs.

Although it is not a strong requirement, the main implementation issue arises with the need to exchange information between client stations and APs. Since the needed information exchange is related to radio measurements, this requirement will be satisfied with the advent of new IEEE standards: IEEE 802.11h and 802.11k

Algorithm 3 *txPowerWill* at AP j

```
Require: $S^j_{rg}$, $SNR_i, SNR_{j,i}\ \forall i \in S^j_{rg}$
Ensure: $Pt_j$
  $end \leftarrow false$
  while !$end$ do
    if $Pt_j + step > PtMax$ then
      $end \leftarrow true$
    end if
    for all $S_i \in S^j_{rg}$ do
      if $(SNR_i < CS_{Th})$ && $(SNR_{j,i} > SNR_i + \Delta SNR + step)$ then
        $end \leftarrow true$
      end if
    end for
    if !end then
      $Pt_j += step$
    end if
  end while
```

Algorithm 4 *txPowerGull* at AP j

```
Require: $S^j_{as}$, $SNR_i, SNR_{j,i}\ \forall i \in S^j_{as}$
Ensure: $Pt_j$
  $end \leftarrow false$
  while !$end$ do
    if $Pt_j - step < PtMin$ then
      $end \leftarrow true$
    end if
    for all $S_i \in S^j_{as}$ do
      $MaxSNR_i = argmax_{k \in A}(SNR_{k,i})$
      if $(SNR_i < CS_{Th})$ && $(SNR_{j,i} - step < MaxSNR_i + \Delta SNR)$ then
        $end \leftarrow true$
      end if
    end for
    if !end then
      $Pt_j -= step$
    end if
  end while
```

Chapter 7

Resource Management in WLANs

In the previous chapters we proposed different radio resource management techniques intended to improve the network performance. Although some implementation issues were already discussed, this chapter provides full details on how those techniques are integrated within a management environment.

Nowadays, most IP data networks are managed by using the popular IETF's Simple Network Management Protocol (SNMP), while telecommunications networks traditionally use the ITU-T's recommended scheme known as Common Management Information Protocol (CMIP). The current IEEE 802.11 specification implies that stations may be managed via SNMP. However, the use of SNMP introduces the following problems: 1. Only a reduced set of stations in the market include SNMP capabilities. 2. The use of secure SNMP protocol (e.g. SNMPv3) requires significant pre-configuration of the station. 3. Management of a station may be required prior to the establishment of an IP connection. There are cases where a device must be managed because it cannot get IP connectivity.

Besides, in [177], the performance of SNMP in radio-link networks is evaluated. According to this study, the key parameter is the response delay, given its highly variability which depends on the channel status, and network load. In the worst case scenario, this delay limits the number of managed devices to 140.

Furthermore, administrators have no choice but to share the ISM band with devices belonging to other networks; that is to say, not all the devices interacting with the network are manageable. All these issues, together with the dynamism of the radio environment, which would require a rapid adaptation regardless of the number of managed devices, are changing the traditional centralized approach towards distributed architectures. Although SNMP allows slightly distributed architectures, supporting the configuration of hierarchies and delegation of tasks among managers of different levels (SNMPv2 and beyond), it still follows a similar centralized paradigm and its efficiency depends on the capacity of the manager entity.

As the capabilities of network devices increases, more attention is paid to distributed systems. The increasing decentralization produces noteworthy benefits [173]: reduces the load on the bottlenecks (managers), both in traffic and in processing load; improves robustness since it reduces dependency on a single point of failure; and enables a faster response to changes since the "managed" devices now have more autonomy to take their own decisions.
The authors in [96] propose an analytical model that quantitatively assesses the efficiency of different management schemes in terms of delays, computational load and security. Centralized systems are outperformed by distributed schemes, and they can only be improved with the use of hierarchies, delegating tasks to lower level managers; the more evident the decentralization is, the better the performance in terms of delay and process load. However, decentralization may be translated into increased signaling traffic.

All these arguments are not intended to promote the death of centralized management. As a matter of fact, the centralized approach is dominant today and will remain so in the future, mainly due to its simplicity and ease of implementation. The recent development of CAPWAP, which is detailed later on this chapter, supports this thesis. For these reasons we explored both ways and came up with two different approaches, which are detailed in this chapter: a centralized scheme intended to leave all control in the network operator's hands, and a distributed approach, based on "smart" devices that cooperate. The first approach, namely UAMN, was published in [243, 244] and consists of a series of WLAN APs interconnected through a mesh network with an additional cellular interface that provides contact with a central manager within the operator's domain. The second, published in [127], involves a set of neighboring APs that are able to cooperate in order to optimize the utilization of a shared resource. As mentioned above, the latter implies a special attention to signaling.

Furthermore, new IEEE 802.11 standards provide interesting features as regards radio resource management. Part of the required signaling could be solved by means of these new functionalities. For this reason, we thought it would be interesting to begin this chapter with the most relevant aspects of these standards: IEEE 802.11h, 802.11k and 802.11v.

7.1 Standardized Radio Resource Management

Since the early adoption of the IEEE 802.11 standard as the most popular WLAN technology, the idea of enabling radio resource management was in the administrators' minds. The first approach, the IEEE 802.11h, was intended to solve some regulatory constraints, motivated by the use of frequency bands that may interfere with radar systems. However, the new functionalities opened a wide range of new possibilities for enhancing spectrum efficiency: 802.11h enables power control and facilitates dynamic channel management.

Later, the IEEE 802.11k provided a new means to exchange information related to radio related measurements. This is indeed an important feature that allows any resource management application to provide better results, since more precise information on the radio channel is available.

Finally, TGv is still working on an amendment that will allow, among other features, the management of client devices from the network side.

7.1.1 IEEE 802.11h: Spectrum Managed

Spectrum management services are a special subset of station services. They are designed to allow the wireless network to react to conditions and change radio settings dynamically. The development of the IEEE 802.11h [14], now included as part of the current 802.11 standard [24], follows an ITU recommendation that was motivated by the European Radiocommunications Office (ERO)[1] demands of minimizing the impact of opening the 5 GHz band, generally used for military systems, to ISM applications [16]. The ETSI regulatory domain required the inclusion of new features not found in the initial 802.11a products brought to market. Two services were defined in 802.11h to help meet those regulatory requirements: Transmit Power Control (TPC) and Dynamic Frequency Selection (DFS). Both features were handled quite well by the HiperLAN 2 specification [7], a technology that was in competition with 802.11a.
In 2004, the FCC opened up the frequencies between 5.470 and 5.725 GHz, provided that DFS and TPC are implemented properly.

7.1.1.1 TPC

TPC allows the transmission power of a STA to be dynamically adjusted. APs will be able to use the TPC operations to advertise the maximum permissible power, and reject associations from clients that do not comply with the local radio regulations. Although designed to satisfy regulatory requirements, transmit power control may have additional benefits. For example, an efficient TPC will hold transmit power to the lowest possible productive level. In this way, clients can use TPC to adjust power of transmitted frames so as to reduce interference and save battery, while at the same time the quality of the transmitted signal is enough to guarantee a correct reception using a “fast” modulation. Power control also helps to simplify the electronics in the AP because all signals will be received at roughly the same level.

Regulatory maximum power may be configured into an AP or STA, or it may be learned from Beacon frames containing *Country* elements. The *Country* element specifies the regulatory maximum power, and the *Power Constraint* element can be used to specify a lower maximum transmission power specific to the network.

[1] http://www.ero.dk

However, the greatest advantage of TPC is that it allows dynamic transmitted power selection on a per-packet basis. For each frame, the receiver may compute the link margin. The link margin is defined as the difference between the received power and the minimum acceptable value (which varies with the modulation in use). It provides a safety margin that guarantees that the minimum received power is met with a high probability in spite of detrimental changes in the propagation path. IEEE 802.11h provides means to make informed changes to transmission power attending the desired link margin. New *Action* frames are defined by the 802.11h to allow a STA to request a transmission report. The receiving STA responds with a TPC Report with the measurements necessary to make uplink and downlink margin estimations.

7.1.1.2 DFS

The second service, DFS, was developed mainly to avoid interfering with some 5 GHz radar systems in use in Europe. Although originally developed to satisfy European regulators, the underlying principles have been required by other regulators as well. To this aim, DFS will periodically test the channel for potential interference from other radio systems (most notably radar systems, but also other ISM applications). Whenever radar signals are detected, the network must switch to another channel to avoid interference.

DFS includes a way for the AP to quieten the channel so that any STA in its cell can search for other radio activity without interference. It also defines a process to reassign the cell's channel on the fly, based on those interference measurements. The availability of such a feature simplifies the implementation of dynamic frequency assignments. But probably the most significant part of DFS is the way it allows requesting and reporting of radio measurements, settling the principles on which the latter IEEE 802.11k is based (see section 7.1.2).

Measurements can be requested and handled by any STA, at any moment. Those on-demand measurements can be particularly useful for any radio resource management. The AP may ask other STAs to make a measurement. Note that those requested reports from stations are likely to come from a variety of different geographical locations, thus allowing the AP to learn about the state of wireless medium beyond its coverage area.
Inquiries for radio information are sent using *Measurement Request* frames. In addition to the polled operation, in which the AP (or other STAs) ask for measurements, it is possible to report statistics spontaneously by sending unsolicited *Measurement Report* frames triggered by certain events. Measurements may be taken either during a quiet period or while the radio is in service.

As mentioned above, quiet periods or quiet intervals are used to perform measurements on the radio channel. In an ad hoc network, the quiet period scheduling is chosen when the network is created, whereas in an infrastructure WLAN this scheduling is completely under the control of the AP. Quite intervals are announced by the AP using the *Quiet* information element present in *Beacon* and *Probe Response* frames.

During a quiet interval all STAs in the BSS must cease their transmissions. This is achieved by setting the STAs' NAV to the length of the quiet period. As explained in section 2.3.2, the NAV is used for virtual carrier sensing to defer transmissions as if the medium were busy. When the quiet period resumes, all stations must contend for access to the radio channel again. There is no preservation of channel access across a quiet period.

Another interesting feature provided by the DFS service is the channel switching mechanism; that is, 802.11h enable a network to move seamlessly to another channel. A channel switch may be produced by either the presence of a radar in the channel in use, or because an overlaying channel assignment mechanism has decided to change the frequency plan (e.g. see chapter 5). The APs (in case of infrastructure WLANs) or a selected STA (in case of IBSS) inform associated STAs of the impending switch by using the *Channel Switch Announcement* information element in management frames, as well as *Action* frames.
As detailed in Section 5.2.3, in the absence of this mechanism, an uncoordinated channel switch by the AP may cause associated STAs to loose connectivity for long periods, before they realize that the BSS has moved to another channel.

7.1.1.3 Action Frames

The IEEE 802.11h amendment added a new management message subtype, *Action* frames. There are five different *Action* frames: *Measurement Request*, *Measurement Report*, *TPC Request*, *TPC Report* and *Channel Switch Announcement*.

Measurement Request/Report a *Measurement Request* frame is used to request that a station make measurements and send the results to the sender through a *Measurement Response*. The requesting STA may specify the moment at which the measurement should start, and its duration. Periodic responses set by a single request are also allowed. Three different types of measurements can be requested:

Basic Measurement for a given channel, the requested STA responds with a series of flags determining whether it has detected a valid frame or not, an OFDM preamble, radar presence, or any unidentified signal source exceeding a given threshold.

Clear Channel Assessment (CCA) describes the fraction of time on which the requested STA's CCA function was set to busy (see section 2.3.2 for more details on the CCA mechanism).

Receive Power Indication (RPI) Histogram is used to report the spread of received power on an interface. The histogram contains eight fields, each of which represents a range of received power. Each field has a value representing the fraction of detected signals that falls into its power range.

TPC Request/Report message exchange enables a STA to estimate uplink and downlink margins. The responding STA reports the amount of energy by which the received *TPC Request* exceeds the minimum acceptable value (i.e. the uplink margin), and also contains the transmit power used to send the report frame. Based on this information, the requesting STA can estimate the path loss of the radio link.

Channel Switch Announcement is used to announce the time at which the network will switch to a new channel, and of course it includes the new channel. The AP may send the *Channel Switch Announcement* frame without performing a backoff, after determining the medium is idle for one PCF Interframe Space (PIFS) period. In this way, the announcement gets a higher priority than any new atomic exchange on the medium.

7.1.1.4 New information in management frames

Besides defining new types of frames, the IEEE 802.11h also added new information elements in previously existing frames.

Beacon and *Probe Response* frames are used by the AP to define the regulatory domain and the power constraints that apply on its cell. The AP also includes a *TPC Report* element that is sent without the corresponding request. In this case, the link margin field is set to zero and must be ignored. Additionally, *Beacons* and *Probe Responses* are used to announce quiet periods and channel switches.

Furthermore, STAs use the *Association* and *Reassociation Request* frames to inform the requested AP of their power capabilities and the list of supported channels. More precisely, the power capability element specifies the minimum and maximum transmit powers with which a STA is capable of transmitting in the current channel.

7.1.2 IEEE 802.11k: Radio Measurements

The work of the TGk group (Radio Resource Measurement) resulted in the IEEE 802.11k amendment [36], released in December 2008. According to their own words, Radio Resource Measurement addresses some of the existing issues in using unlicensed radio environments to meet the requirements of emerging technologies. IEEE 802.11k provides knowledge about the radio environment to improve performance and reliability. This improved knowledge is achieved by means of standardized radio measurements. The scope of TGk was therefore to define those radio measurements and to provide mechanisms to higher layers for radio and network statistics. In other words, the proposed Radio Resource Measurement approach is to extend the capability, reliability, and maintainability of WLANs through measurements and provide that information to upper layers in the communications stack.

The measurement of radio parameters in a WLAN allows an automatic and dynamic adjustment to the wireless environment at application level. In 802.11k enabled WLANs radio measurements are not only carried out locally, the standard allows any STA to send measurement requests to other STAs. The measurement results are made available to upper layers for any purpose, making them especially useful for RRM. That is to say, 802.11k is responsible for providing useful information on various radio parameters, any RRM mechanism definition is out of its scope. In summary, the new services that TGk added to the 802.11 provide the following:

- Ability to conduct radio measurements in supported channels.
- Ability to make requests and report radio measurements via the wireless interface.
- An interface for upper layer applications to retrieve radio measurements, whether it be using MAC Sublayer Management Entity (MLME) primitives and/or MIB access.
- Provision of information about nearby APs.

In order to provide the ability to request or report a collection of radio measurements, IEEE 802.11k inherits mechanisms that were previously defined by the TGh (see section 7.1.1 for more details). For example, 802.11k uses 802.11h's *Measurement Request/Report Action* frames. More precisely, it extends the family of *Action* frames with the definition of two new categories: *Radio Measurement*, and *Public Action* frames. Furthermore, the 802.11k amendment extends previously existing management frames to include diverse radio measurements.

The implementation of the IEEE 802.11k specifications involves extending the SNMP MIB to provide access to new objects describing radio link parameters.

7.1.2.1 New responsibilities for the old STAs

As mentioned before, with 802.11k the STAs gain the ability to perform radio measurements on the channel. The resulting information can be made available to either local upper-layer entities or to other neighboring STAs after a request.

Radio measurements can take place in the working channel, so that ongoing data sessions are not disrupted. These kinds of measurements are called concurrent measurements. However, radio measurements can also be carried out in different channels (called dedicated measurements). In this case, the STA must save state information in order to maintain the current association with a BSS or an IBSS before it switches to the channel where the measurements take place. For this reason, dedicated measurements should be requested less often and for shorter durations.

The standard allows the duration of the measurement to be defined in the request message. This duration can be mandatory or either a “recommended” time that the measuring STA will try to meet as far as possible. The requested STA will try to start the measurement immediately after the request, although in some cases it is advisable to apply a random delay prior to the measurement or between consecutive measurements in order to prevent flooding of responses after a multicast/broadcast request.

Requests can be sent either to individual STAs or to a group of STAs (multicast/broadcast) in the same BSS or IBSS. In an infrastructure-based BSS, non-AP STAs are only allowed to send measurement requests to their AP. When a STA has more than one pending measurement, it will give higher priority to unicast requests, given that multicast measurement requests might already have received a response from other STAs. Similarly, upon refusal to perform any requested measurement (e.g. the STA is unable to perform the measurement), the requested STA must inform the requesting STA, but only in the case of unicast requests.

As with 802.11h, STAs are also allowed to send measurement reports asynchronously (without a previous explicit request). These unsolicited reports are triggered when a particular parameter under measurement exceeds a given threshold. The thresholds that trigger an automatic report are set out in an initial request and can be modified in subsequent requests. The automatic reporting is stopped when the reporting STA moves to another BSS or after an explicit request.

7.1.2.2 Action frames

The IEEE 802.11k amendment added new management messages within the *Action* subtype intended to request and report specific radio measurements. There are two new categories of *Action* frames: *Radio Measurement*, and *Public*. In the first group, there are *Radio Measurement Request/Report*, *Link Measurement Request/Report* and *Neighbor Report Request/Response*. In the second group, a new frame is included: *Measurement Pilot*.

Measurement Pilot frames are compact *Action* frames that are transmitted pseudo-periodically by an AP to the broadcast address, similarly to *Beacon* frames. The difference between *Measurement Pilots* and *Beacon* frames is that the latter provide much more information and are transmitted less frequently. The purpose of the *Measurement Pilot* frame is to help STAs to quickly detect APs in the area without an active scan (see section 2.2.2.1 for more details on the 802.11 scanning process). The AP includes very little information on a *Measurement Pilot* frame although, as in other frames, the standard allows additional optional and vendor-specific information to be included.

Neighbor Request/Report are used to facilitate handover decisions taken by client STAs. A non-AP STA may request a list of APs in the area by means of a *Neighbor*

Request. This request is sent to the AP to which the STA is currently associated. In response, the AP sends the list of neighboring APs (members of the same ESS) stored in the newly defined *dot11RRMNeighborReportTable* of the MIB. The standard does not define how this table should be updated. It may have been set by hand by an administrator, in which case it may contain stale information.
For each AP in the list, the reporting AP includes:

- BSSID
- AP's capabilities (e.g. 11e QoS enabled, 11h spectrum management, etc.)
- AP reachability (from the requesting STA) for a possible pre-authentication
- AP's working channel
- Supported PHYs
- Other optional and vendor-specific information

Link Measurement Request/Report allows a STA to assess propagation losses in the path to another STA. The STA sending a *Link Margin Request* includes the power (in dBm) used for transmitting the request message. Then, a station receiving the request shall include in its response the power and SNR at which the request message was received. It also includes information on the transmission and reception antennas and an 802.11h's TPC Report element (see section 7.1.1.3). Briefly, a request/reply exchange of this kind allows two STAs to assess the minimum transmitted power required to maintain a bidirectional communication.

Radio Measurement Request/Report is the generic message exchange used to request and report one or more radio measurements. The list of available measurements comprises a wide collection of parameters that provide comprehensive information on the channel status. In addition to the measurements described in the following paragraphs, *Radio Measurement Action* frames can be used to obtain 802.11h's *Basic*, *CCA* and *RPI histogram* measurements.

Beacon Report is used to obtain information on the signal level of received *Beacon*, *Measurement Pilot* or *Probe Response* frames. These frames are collected from one or more BSSs found in one or more channels. This report may assist in identifying potential AP candidates for an eventual handover. Sending a report of this type may be conditioned upon receipt of power above or below a given threshold. Among other information, for each *Beacon* detected, the *Beacon Report* includes the following:

- Received power (dBm) and SNR (dB)
- Frequency channel
- PHY type
- The body, or a portion of the frame (as explicitly detailed in the request)

A Beacon Report may be created in three different ways: *Passive*, *Active* or *Table-based*. In the first case, the requested STA passively listens for *Beacon* and *Probe Response* frames. In the second, the requested STA forces the transmission of *Probe Response* frames by broadcasting *Probe Request* messages. In the latter case, the STA sends stored data obtained in previous measurements.

Frame Report enables different measurements based on any frame received from a particular STA or from a group of STAs. However, the standard only defines the *Frame Count Report*. The STA that performs such a measurement observes the traffic on a given channel and makes a summary for each different transmitter detected during the measurement. Similar to the *Beacon Report*, for each transmitter, the requested STA informs of received power, SNR, etc. A *Frame Count Report* includes the number of data or management frames received from a given STA.

Channel Load Report provides the load metric of a given channel based on the measurement of busy time (see Section 6.1.4.2). More precisely, channel busy time is defined as the proportion of the time during which either the physical carrier sense, the virtual carrier sense (NAV) or both indicate that the channel is busy. This measurement is similar, although not identical, to the *CCA* report added by the IEEE 802.11h. The latter does not use NAV information.

Noise Histogram Report presents a measurement of received power detected while the carrier sense mechanism indicates a free medium.
The histogram consists of ten different levels (between -92 and -55 dBm) and is drawn by taking into account the portion of time in which noise is detected in each of the ten levels. This measurement is similar in conception to the 802.11h's *RPI Histogram*, but besides the power densities, the *Noise Histogram Report* shall include the average noise plus interference power on the measured channel at the antenna connector during the measurement duration.

STA Statistics Report is not actually a radio measurement per se. Upon receipt of a *STA Statistics Request*, the requested STA sends information stored in different MIB objects. Some of them belong to the current definition of the 802.11 standard (including the 802.11e amendment): *dot11Counters* (OID: .1.2.840.10036.2.2), *dot11MacStatistics* (OID: .1.2.840.10036.5.1.4) and 11e's *dot11QosCountersTable* (OID: .1.2.840.10036.2.6). In addition to the mentioned statistics, the report may include measured MAD values (see Section 6.1.4.1).

Location Configuration Information Report is used to provide presence information. Requests can be either *Local*, or *Remote*. Using the first type of request, a STA seeks to obtain its own location (i.e. requesting STA says - where am I?). The second request is used to obtain another STA location (i.e. - where are you?).

These reports include Latitude, Longitude, Altitude, and optional Azimuth information of a given STA, following the format described in RFC3825.

Transmit Stream/Category Report consists of measurements of the quality provided to a given Traffic Stream (TS) or Traffic Category (TC). Obviously, this measurement only applies to QoS-capable STAs. The measurement is performed on the traffic between the requested STA and a target STA (specified in the request). The quality is measured in terms of delay, transmission attempts, frame loss, etc.

7.1.2.3 New information in management frames

Besides the TPC element, already introduced by 802.11h, an IEEE 802.11k enabled AP includes the following information in *Beacons* and *Probe Responses*:

- *AP Channel Report*: list of channels where an AP is likely to be found (according to the MIB element *dot11APChannelReportTable*).
- *BSS Average Access Delay*: MAD load metric, as defined in section 6.1.4.1.
- *BSS AC Access Delay*: in QAPs, MAD values measured for the four different ACs (Best Effort, Background, Video and Audio).
- *BSS Available Admission Capacity*: in QAPs, the remaining amount of medium time available via explicit admission control (i.e. the original definition of AAC).
- *Antenna Information*: information on the antenna used to send that frame.

It is worth mentioning that early versions of the IEEE 802.11k standard also added the *BSS Load* element in *Beacon* and *Probe Responses*. However, TGe went ahead and included this information element on their amendment (released three years before). The *BSS Load* element, which has been previously defined (see section 6.1.4), provides interesting information for the mechanisms discussed in Chapter 6. An in-depth analysis of those parameters can be found in [268].

Another interesting feature added previously by the 802.11d and included in the current standard [24] is that a STA can request any information element in a *Probe Request* frame. The AP will then respond with a *Probe Response* carrying the requested information, which may include any new element (i.e. radio measurement) introduced by the IEEE 802.11k.

No new changes are made to *Association Request* frames with respect to the 802.11h definition, which included information about the minimum and maximum transmit powers with which a STA is capable of transmitting in the current channel. In contrast, the *Association Response* now bears information about the power and SNR at which the request was received. The uplink margin is thereby assessed by the requesting STA without the need of a *Neighbor Request/Report* exchange.

7.1.3 IEEE 802.11v: Network management

This amendment is still in the early stages of development. A first draft emerged in July 2007, but the final 802.11v standard is not expected until July 2010. The current version of the Draft document is v3.01 [37] (at the time of writing), but like other IEEE 802.11 working documents, it is not publicly available. For this reason, the contents of this section may be outdated or may not reflect the actual definition of all IEEE 802.11v's processes and mechanisms.

As stated by the TGv, the IEEE 802.11v amendment defines mechanisms for wireless network management of non-AP STAs, including BSS transition management, co-located interference reporting, diagnostic and event reporting, a traffic filtering service, power saving enhancements and presence. To this end, IEEE 802.11v defines a new set of *Action* frames: *Wireless Network Management Action* frames.

7.1.3.1 Power saving and traffic filtering

One of the most publicized new features added by 802.11v is the improved Power Saving Mode (PSM). In the current IEEE 802.11 standard, a STA using PSM only leaves the dormant state to send frames or to wake up periodically in order to receive frames. Unicast packets directed to a STA in PSM are buffered in the AP; each STA has its own time interval to wake up and see if there is anything in its buffer. However this mechanism showed some limitations. The limitations of PSM motivated the development of several improvements, some of which are gathered under the Wi-Fi Multimedia Power Save certification (WMM-PSTM)[2]. WMM-PSTM includes the Unscheduled Automatic Power Save Delivery (U-APSD) to achieve a reduction in latency and overhead in unicast frame delivery. However, these improvements do not change the legacy PSM multicast/broadcast frame delivery mechanism. This mechanism unnecessarily forces STAs to wake from doze state in order to deliver multicast/broadcast frames even though the station is not interested in them. The upcoming 802.11v will allow STAs to define the multicast traffic they are interested in, and therefore they will only wake up when necessary. More precisely, a STA may send a *Traffic Filtering Service Request Action* frame to the AP to request the specified traffic filtering. The request includes an octet string that is compared to the frame content so as to decide whether this downlink frame is delivered or not. Moreover 802.11v includes the definition of an ARP Proxy mechanism that filters Address Resolution Protocol (ARP) requests (sent to the broadcast address). Those broadcast requests wake up all dormant STAs regardless of the intended destination.

7.1.3.2 Measurement, event and diagnostic reporting

IEEE 802.11k radio measurements are extended by the TGv to add the new *Multicast Diagnostic Report*. This mechanism allows an AP to receive statistics of broadcast/multicast

[2] http://www.wi-fi.org/knowledge_center/wmm

traffic at associated STAs. Any STA accepting such a request shall count the number of received MSDUs with the specified multicast address during the requested duration. Another interesting measurement is provided by the Co-located Interference reporting. It allows a requesting STA to receive information concerning the co-located interference being experienced by another STA on the operating channel. The requested STA sends a *Co-located Interference Response* frame to the requesting STA if it detects that a co-located radio or another co-located interferer is causing performance degradation to its WLAN receiver. This report may include the periodicity, level of interference, accuracy of the reported interference level, interference center frequency and interference bandwidth.

Besides radio measurements, 802.11v also enables the reporting of different MAC level events through the *Event Request/Report* exchange. By using this feature, an AP is able to request its associated STAs information about recent or past transition events (i.e. handovers), RSNA events regarding authentication parameters, etc. Finally, a STA experiencing network problems may be queried at any time for its current set of vendor-specific syslog messages (in human readable form).

An AP can also use *Diagnostic Request/Report* protocol as a means of diagnosing and debugging complex network issues. For example, an AP may use *Diagnostic Requests* to discover the range of capabilities available to an associated STA or to discover various important operating settings currently in use by an associated STA. These capabilities or settings include available and current data rates, supported channels, the SSID, antenna type and gain, and even the firmware version or the manufacturer, model and serial number.
There are other *Diagnostic Requests* that involve some actions carried out by the requested STA. For example an *Authentication Diagnostic* or an *Association Diagnostic* is used to determine whether a STA is able to perform an IEEE 802.11 authentication (or association) with a designated BSS. This requires the authentication or association processes to be completed by the STA with a specific AP.

7.1.3.3 Presence procedures

An attractive new capability added by 802.11v is the requesting and reporting of location-related information. Regarding this feature, we have to note that IEEE 802.11v is only intended to standardize the way location information is exchanged or advertised, and it does not automatically add location capabilities to WLAN devices. STAs supporting location services need additional hardware (e.g. a GPS) or an enabled interface to an external location information source. However, the exchange of *Radio Information* (antenna gain, transmit power, etc.) together with *Timing Measurement* enables distance estimations between neighboring STAs based on either received power measurements or round trip time calculations. Additionally, STAs can inform about their mobility (e.g. stationary, in motion and the speed).

The *Presence Request/Report* procedure allows any STA to request either its own location information or the remote STA's location, from a peer STA that supports location services. Furthermore, any STA may periodically advertise its presence without a previous request since presence information may be included in *Beacons* and *Probe Responses*, in the case of APs, or in (*Re*)*Association Requests* and *Presence Requests/Responses*, in the case of non-AP STAs.

7.1.3.4 Association control

There is still another new feature that is of special interest for load balancing schemes, as explained in Chapter 6. Recall that associations between STAs and an APs are decided and handled by the STAs based on signal strength measurements of received *Beacon* (and *Probe Response*) frames. The *BSS Transition Management* mechanism defined by the IEEE 802.11v provides means to perform network-driven association management. That is to say, the network is now able to decide the AP to which a given STA must be associated, based on load balancing criteria. The *BSS Transition Management Request* is sent from an AP to ask a STA to move to another (less loaded) AP. The request may include a list of candidate APs in the form of 802.11k's *Neighbor Reports* (see section 7.1.2). Those *Neighbor Reports* are extended by IEEE 802.11v to include cell load information, enabling client-driven load balancing. What is more, 802.11v not only enables network-directed load balancing, but also promotes and improves the traditional client-driven association scheme since any STA may request information from the AP through a *BSS Transition Management Query* to determine the timeliness of an intended handover.

7.2 Centralized management architectures

Centralized WLAN architectures indeed simplify the deployment of large-scale networks by enabling network-wide monitoring and by improving configurability.

For these reasons, centralized management has been traditionally used for both the wired and the wireless worlds, since the advent of the SNMP protocol. As a matter of fact, there are several commercial solutions on the market for the centralized monitoring and configuration of large-scale WLANs. For example, Aruba Networks' AirWave [28] or Enterasys Wireless Management Suite [40] offer a management environment for multi-vendor WLAN devices based on SNMP. Besides traffic monitoring, user location and tracking, and security alarms, these solutions are claimed to optimize APs' radio parameters. However, no hints are given on what exactly they optimize or how it is done.

In [169], the authors provided the basis for the development of an SNMP-based management application, specifically designed for 802.11 WLANs. In [298], SNMP is used to implement a network-driven load balancing mechanism: clients request an association to a central manager, which maintains load statistics from the controlled APs by means of SNMP.

Strongly related to SNMP, other proposals use the newly CAPWAP architecture for radio resource optimizations. Cisco, precursor of the IETF CAPWAP initiative provides the Wireless Control System [29] to manage a collection of light-weight APs (although it also supports a limited set of standalone "thick" APs). In [200] and [66], simple wireless network resource management functions are proposed to investigate the applicability of CAPWAP. In [44], a similar approach is proposed, based on the same idea of managing "thin" APs.

In the following sections, we describe the basic operation of SNMP, since it is still the "de facto" standard for (centralized) network management. Besides, given its recentness and the fact that it presents an alternative to our vision of "intelligent" devices, we considered that an introduction to CAPWAP was also worth providing. Finally, our centralized approach is explained. We have to note that the UAMN is not intended to substitute or compete against SNMP or CAPWAP, since it is not a management protocol. UAMN is a novel architecture which includes a management framework for APs that makes use of 802.11 and cellular interfaces.

7.2.1 SNMP

Simple Network Management Protocol (SNMP) is a component of the Internet Protocol Suite defined by the Internet Engineering Task Force (IETF) in several RFCs. It is considered nowadays as the *de facto* standard for management and administration of IP networks. The most recent version of SNMP is known as SNMPv3 [87]. The set of standards comprised within the SNMP management scheme include the definition of an application layer protocol, a database schema, and a set of data objects. These objects are in fact configuration parameters and statistics that are exposed in the form of variables on the managed systems. These variables can then be queried (and sometimes set) by managing applications.

SNMP defines a centralized architecture [153] (see Fig. 7.1), where one or more network managers handle several network devices through their SNMP agents.
The network manager can be described as an application software used to monitor and control managed devices. This entity allows the centralized management of the network and thus provides the bulk of the processing and memory resources required for those processes. The SNMP agents are the entities that actually interact directly with the devices. An SNMP agent is a software module running in every managed device and is in charge of managing local information and configuration parameters according to the network manager's indications. Managed devices also contain managed objects. As explained before, those objects can be used to represent hardware modules, to define configuration parameters or to store statistics and other data. The objects, in the form of variables, are made accessible via SNMP by the agents and are organized in hierarchies.

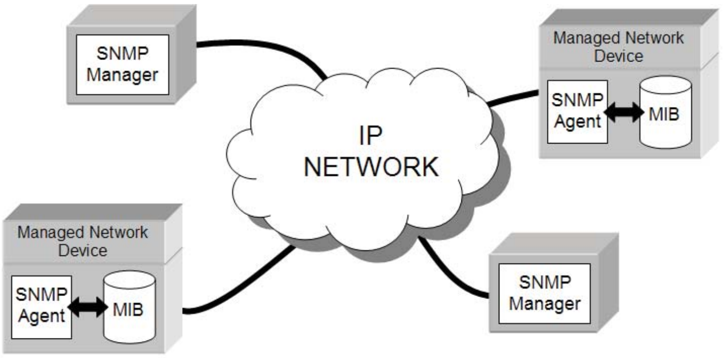

Figure 7.1: Architecture for SNMP frameworks

These hierarchies, and other metadata (such as type and description of the variable), are described by Management Information Bases (MIBs).

Communication is needed between agents and managers to access the objects stored in the MIB. SNMP uses UDP as the transport protocol, and therefore it assumes that this communication is not connection oriented, so there is no guarantee that an SNMP message has been successfully delivered. This makes the network manager take the responsibility for all reliability issues. SNMP dialogue follows a request-response scheme, but it also defines an event-driven mechanism. SNMP Agents are allowed to report an alert or other asynchronous event to the manager, even though it has not been previously requested. Such events can be of different nature, e.g. a given statistic exceeds a given threshold, the managed device has just recovered after a power-off, etc.. This type of asynchronous message is called *trap* (in SNMPv1 [88]) or *notification* (in SNMPv3).
The event-driven communication between agents and managers puts forward interesting applications. For example, in [132] we presented a centralized location system for IEEE 802.11 WLANs that used SNMP traps. Every time an AP completes the association of a new STA, it sends a trap reporting the association. Then, the manager receiving those traps knows the AP (and hence the room) to which a STA is associated.

Following this Manager/Agent paradigm, managed objects must be accessible so that their information can be either retrieved or modified. The Structure of Management Information (SMI) is defined by the IETF in RFCs 1155 [261], and 1212 [262]. The SMI establishes how the information is organized, named and described in order to allow a logical access.
Each object has a unique object identifier (OID) consisting of a series of numbers separated by periods that depend on the object's position within the tree defined by the SMI (see Fig. 7.2); for example, the standard group objects related to TCP/IP stack can be found under the OID group .1.3.6.1. According to the SMI, the structure of such objects is described according to ITU's Abstract Syntax Notation One (ASN.1). The mentioned notation describes data structures for representing, encoding, transmitting, and decoding data. It provides a set of formal rules for describing the structure of objects that are

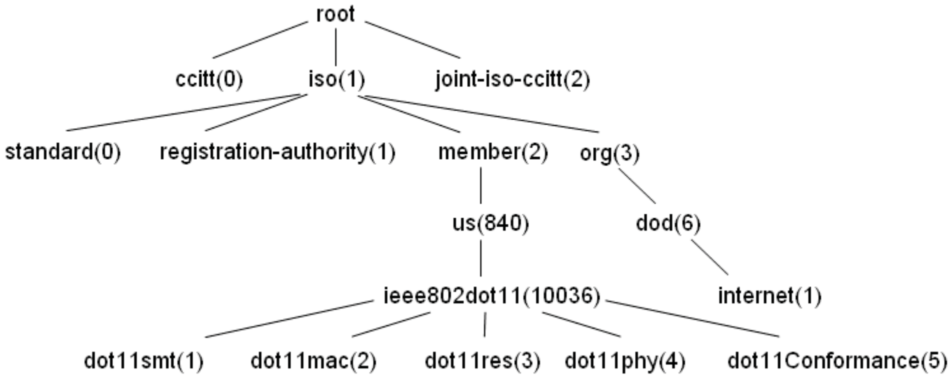

Figure 7.2: SMI tree for ieee802dot11 group

independent of machine-specific encoding techniques. A MIB is then a collection of definitions that specify the properties of objects in a managed device. The device maintains a database of values for each object defined in the MIB.
The base MIB, which defines the most important objects used to define the TCP/IP stack, is known as MIB-II and is defined in RFC 1213 [220]. Of course, each different type of device may contain a set of particular objects, in addition to the common objects of the MIB-II. There are a number of RFCs that define specific MIBs for certain devices.

The IEEE 802.11 contains an extensive collection of management functions. The complexity of the additional variables added by the 802.11 has been organized in a specific MIB whose formal description is detailed in Annex D of the IEEE 802.11 specification [24]. The different amendments of the standard update this structure by modifying or adding new parameters.
The specific group of objects devoted to specific 802.11 layer 1 and layer 2 parameters is located under the OID .1.2.840.10036 (see Fig. 7.2). This MIB is described and discussed in detail in [142]. The four main groups are briefly described below.

1. ***dot11smt***: contains objects related to the station management and the local configuration. Includes global configuration parameters that are not part of the MAC itself (e.g. ESSID, multi-domain capability, infrastructure mode, ad-hoc, etc.), authentication, and privacy, and provide a means for automated notification of significant events.

2. ***dot11mac***: this group provides access to objects that allow tuning and monitoring MAC layer parameters. It contains useful statistics (transmitted fragments, FCS errors, etc.). It also allows the multicast processing at layer 2 to be configured.

3. ***dot11res***: contains objects that describe available resources. The standard does not define any accessible objects it is rather intended to provide room for MIB extensions defined for a specific model, or a particular manufacturer.

4. ***dot11phy***: This group contains information about different physical layers (DSSS, FHSS, OFDM, and Infra-Red). The information provided for DSSS and OFDM PHY is very limited; it includes information on the frequency channel, CCA mode, or the use of short or long preamble in HR-DSSS. It also provides support for antenna diversity and setting of transmission power (in 8 levels). Also there is space for specific objects of certain products and manufacturers.

In brief, even after extending the IEEE 802.11 MIB with the amendments defined by the TGh [14], there is little information related to radio resources. In addition to this, there is a lack of uniformity in the provision of such parameters among different manufacturers. As mentioned, the IEEE 802.11 left room for private MIBs to fill dot11res or dot11phy groups, and in consequence, a given parameter may have different names and be in different units, depending on the device manufacturer. Therefore, SNMP is not powerful enough by itself to enable an improved radio resource management, unless a comprehensive collection of RRM manageable objects and statistics is standardized and implemented by manufacturers.

7.2.2 CAPWAP

The Control and Provisioning of Wireless Access Points (CAPWAP) IETF group[3] is on the point of finalizing (at the time of writing) the standardization process of a new protocol and a set of mechanisms intended to solve the problems of implementing large-scale WLANs.
The problems introduced by the large-scale deployment of 802.11 WLANs in enterprise networks were analyzed in RFC 3990 [238]. In that document, four basic issues are put forward: first, that each AP is a new networking device requiring management, monitoring and control; the second point states that distributing and maintaining a consistent configuration among a large number of APs (may include a great variety of manufacturers and vendors) is a difficult task, because, although part of this configuration is unchanged in the long run (IP address, ESSID, etc.), another part is completely dynamic and requires a constant attention hindering the provision of a unique and consistent configuration; thirdly, coping with the dynamic nature of the wireless medium requires a coordinated control for maximizing network performance; finally, there are evident security issues, so there is also a need to provide safe access to the network and prevent installation of unauthorized APs.

CAPWAP inherits many concepts from its predecessor, Cisco's Light Weight Access Point Protocol (LWAPP) [82], maintaining its philosophy. The idea is not only to maintain a centralized management of a series of "light" APs, but also to move part of their "intelligence" to a central controller. According to the RFC 4118 [296], an IEEE 802.11 access network is divided into two device categories: Wireless Termination Points (WTPs),

[3] http://www.ietf.org/html.charters/capwap-charter.html

which in fact represent those “light”’ APs, and Access Controllers (ACs). An AC centralizes the management of several WTPs. The functionality of the AC is not limited to the provision of network configuration; it can also manage WTP firmware loading, authentication and radio resource allocation through RF monitoring and setup, which provides means to optimize global network performance. The goals of this architecture explicitly stated in the current CAPWAP specifications are threefold:

- Centralize authentication and policy enforcement functions for a wireless network and, in some cases, bridging and encryption of user data.
- Move processing away from the WTPs, leaving there only time critical functions and protecting the most critical network parameters in the remote AC.
- Provide a generic encapsulation and transport mechanism for control and management messages, enabling the operation of CAPWAP regardless of the wireless technology.

RFC 4118 proposes three different network architectures, depending on the centralization level of the control operations. The architectures are called *Split MAC*, *Local MAC* and *Remote MAC*. The CAPWAP specifications do not assume any specific wireless technology. However, the RFC 5416 [80] defines the binding for IEEE 802.11 WLANs. Specifically, RFC 5416 details two architectures (*Split MAC* and *Local MAC*) for the case of 802.11 networks. In both cases, the services with strict timing requirements (e.g. beacon generation and probe responses) are kept in the WTP. In a *Split MAC* architecture, Distribution and Integration services (see section 2.2 for a description of these services) reside on the AC, whereas all MAC functionalities are entirely left to the WTP in the *Local MAC* architecture. In the latter case, the AC is responsible for WTP configurations and access policy definitions. That is to say, following a *Split MAC* architecture, the WTP is only devoted to the tunneling of data and management frames (except beacon, probe response and power management frames) to and from the AC.

These advanced control functions require a frequent exchange of management messages between the AC and the managed WTPs. As defined in RFC 5415 [81], CAPWAP control and data messages are sent using UDP, and are secured using Datagram Transport Layer Security (DTLS). Resilience is achieved using a request/response scheme, where timeouts cause retransmissions when a response does not follow a certain request.

The CAPWAP protocol also defines a discovery protocol that enables automatic association of WTPs to the AC. Whenever a new WTP is initiated, it sends a *Discovery Request* message and waits for an AC *Discovery Response*. Once the AC is selected, a secure DTLS session is established and both devices exchange their configuration and capabilities.

At the end of configuration phase, the WTP is ready to send and receive CAPWAP messages. Data frames collected by the new WTP are then forwarded to the AC using the

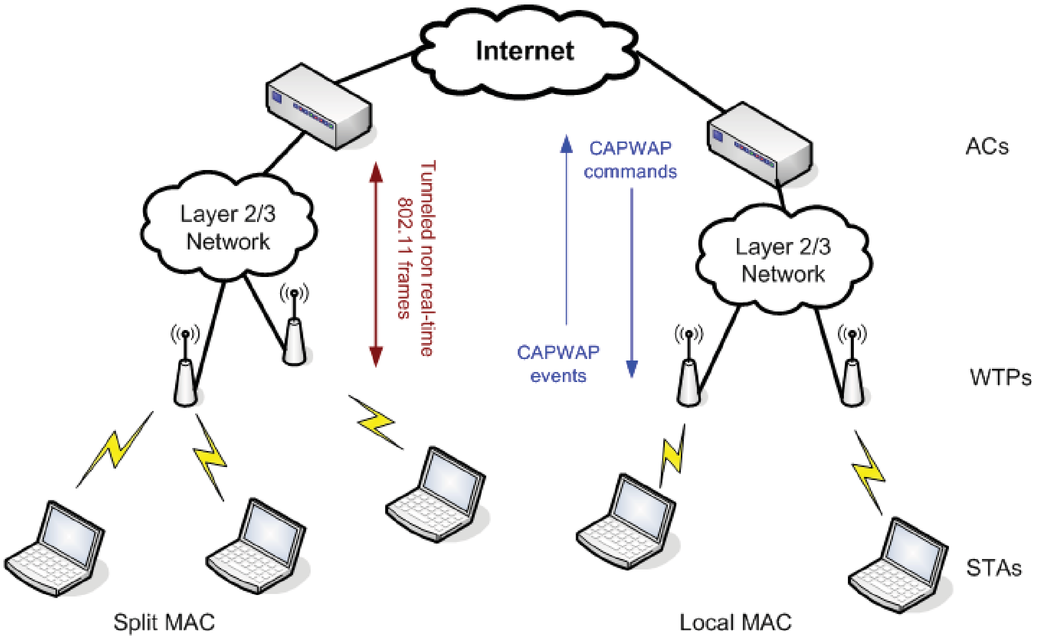

Figure 7.3: CAPWAP network architecture

CAPWAP data message encapsulation rules: IEEE 802.11 header (excluding FCS) and payload are encapsulated and sent to the AC. According to the IEEE 802.11 binding [80], wireless frames encapsulated by the WTP may also include a CAPWAP optional header with information on the RSSI, SNR and data rate used by the sending wireless STA. When the frame is encapsulated by the AC, the CAPWAP header may include information on the WLAN identification to be used when sending the frame to the wireless medium.

Besides data exchange, WTPs and the AC exchange control frames for different management purposes. These control messages can be triggered as a result of either a manual configuration update or automatically generated to dynamically adapt WTP configuration. Control messages allow the AC not only to configure the WTP but also to modify station session state on a WTP (including QoS specifications). The protocol also specifies messages that can be sent asynchronously by the WTPs to notify the AC of certain events (similar to the SNMP traps). Nevertheless, all control messages intended to manage radio resource related parameters are in fact used to read or set ieee802dot11 MIB objects as defined in [24], and therefore the CAPWAP protocol does not provide an added value in this regard. For example *Direct Sequence* message is used in CAPWAP to set the values of the *dot11PhyDSSSTable* (OID .1.2.840.10036.4.5), which allows the frequency channel to be set in a DSSS PHY; the same operation in an OFDM PHY is achieved by means of the CAPWAP's *OFDM Control*, used to set the values of the *dot11PhyOFDMTable* (OID .1.2.840.10036.4.11). Transmission power of a WTP's radio is set through the CAPWAP's *Tx Power*, used to update the value of *dot11CurrentTxPowerLevel* (OID 1.2.840.10036.4.3.1.10).

It is also worth mentioning that CAPWAP extends the WTPs' SMI by adding two new MIBs: the CAPWAP Protocol Base MIB [266] and CAPWAP Protocol Binding MIB for IEEE 802.11 [267]. However, these new MIBs only add CAPWAP's specific timers and variables.

7.2.3 UAMN

Wireless Mesh Networks (WMNs) have been used by the army or in emergency situations due to their robustness and easy deployment. Furthermore, community networks use mesh concepts to skip ownership and the operator's control. These examples do not envisage any commercial exploitation. In [243, 244], our research group studied possible applications of cellular networks in overcoming the limitations of mesh networks, which mainly include security and QoS issues. The project, set up between UPC, Vodafone and Swisscom defines a UMTS/GPRS assisted mesh network (UAMN).

As mentioned earlier, this proposal is not intended to substitute SNMP or CAPWAP, since it covers a wider scope. In fact, UAMN could use SNMP or CAPWAP to remotely configure some mesh node parameters, but on the other hand, none of them cover all the requirements of the whole system.

7.2.3.1 Architecture

The WLAN network consists of several devices that are able to work as APs and as mesh nodes at the same time. To do so, these devices use two different air interfaces: standard IEEE 802.11b/g for the AP interface (provides access to client STAs), and IEEE 802.11a for the backhaul (participates in the mesh network). The mesh network carries traffic between attached users or either between WLAN users and other users/servers in the Internet. This functionality requires the presence of gateway nodes, which are mesh nodes with an additional interface to the cellular network (UMTS or GPRS), or to the wired network (Ethernet, xDSL). Fig. 7.4 shows an example of the architecture of such a network.

If more than one gateway is present, not only reliability and robustness is increased, but also more bandwidth may be offered to the WLAN users. However, if gateways' WAN connection to Internet is built upon cellular links, they become the network bottleneck. In order to mitigate this problem, gateway nodes offer traffic optimization functionalities (e.g., load balancing between gateways and proxy services) which reduce the traffic through the UMTS/GPRS interface.

The UMTS/GPRS network offers a wide coverage and a trusted relation with the user. These two features allow the main drawbacks of mesh networks to be solved. The UMTS/GPRS interface can be used to connect different parts of the same mesh network, manage nodes and monitor the performance of the deployed mesh nodes.

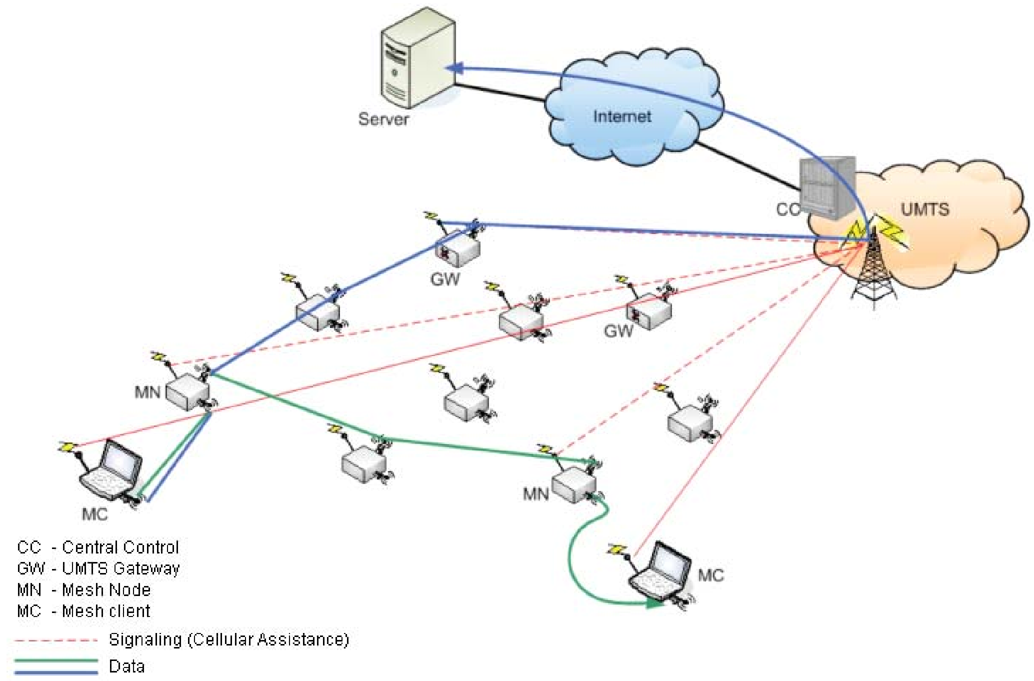

Figure 7.4: UAMN network architecture

User devices and mesh nodes may have a cellular connection in addition to the WLAN interface. This is not a hard requirement on the design but, if present, it eases the implementation of a centralized management scheme. The cellular connection is used by the user and the mesh nodes for signaling, and as mentioned earlier, it may also be used by the gateway nodes for data transmission. The term “network signaling” is used to describe all the messages involved in the user authentication process, accounting and management mechanisms. On the other hand, routing signaling information is only exchanged through the wireless mesh interfaces.

Our aim in this section is to detail the mechanisms related to radio resource management. Nevertheless, this architecture allows us to cope with other requirements, as briefly discussed in the following paragraphs.

Security : as the mesh network can be seen as an extension of the UMTS/GPRS network, the UAMN should offer network authentication and privacy to the user and user authentication and authorization to the network operator. Besides, mesh nodes must be authenticated, and routing information should be sufficiently protected in order to prevent malicious nodes from introducing erroneous information about routes. Additionally, the network should provide self security mechanisms that react upon detecting an attack (e.g., changing functional parameters or avoiding certain routes).

Mobility : the mesh network should be transparent to the user, who will see its connection to the UAMN as a standard WLAN ESS. Therefore, the user must be able to move from one AP to another, ignoring that those APs are in fact, mesh nodes. However, client STAs need to support the security features listed above. For example, the user is only authenticated once, through the cellular interface, so no new re-authentication is needed when the client device carries out a handoff between two mesh nodes. Furthermore, APs (i.e. mesh nodes) are also mobile. This mobility is supported by ad-hoc routing.

Quality of Service : for the same reasoning applied to the security requirements, the UAMN should provide QoS in a similar way as it is provided by the UMTS/GPRS network, but with the additional problem brought by the node mobility and the usage of an unlicensed band. The QoS mechanisms are implemented in mesh nodes and gateways, where the different traffic flows should be transmitted using different paths and the appropriate link quality metrics used in the routing protocol. In order to distribute the traffic efficiently between gateways and perform an adequate admission control, nodes should periodically report (through the cellular interface) information, such as load conditions and link capacity estimations.

Self-configuration : the main advantages of the mesh network are an easy deployment and self-healing capabilities. A node acquires its operational parameters through a dialogue with a central server, but they also learn from the information provided by neighboring nodes. The self-configuring abilities include the frequency assignment mechanism detailed in Section 5.2, load balancing between gateways and mesh path selection.

Protocol Optimization : as mentioned, the mesh network is connected to the Internet using gateway nodes. The GPRS or even the UMTS interface may be the source of some undesired effects on “chatty” protocols, such as Hypertext Transfer Protocol (HTTP), due to long Round Trip Times (RTTs). This behavior degrades the performance of the transmission. This problem can be mitigated with the utilization of protocols that are optimized through the radio interface. In the proposed architecture this functionality is located on gateway nodes, thus requiring minor changes on user terminals.

7.2.3.2 Resource management within UAMN

Two mechanisms presented in this thesis were included in the UAMN set of functionalities: channel assignment (c.f. 5.2) and client-driven load balancing (c.f. 6.2.2). Frequency management requires signaling between APs and the central manager, whereas load balancing involves a signaling exchange between clients and the manager. As detailed in the architecture sub-section, this signaling uses UMTS/GPRS cellular interfaces.

7.2.3.2.1 Load balancing When a new user's device is initiated, it first contacts the central control (CC) through the cellular interface and provides the identities, SINR information and working channel of all APs within reach of its WLAN interface. The central control is then responsible for selecting the best candidate. Note that with all this information the user:

- Provides a list of all APs from other networks that may interfere with that node. This list helps the CC to perform a better frequency allocation algorithm.
- Provides a list of APs that belong to the operator (i.e. the mesh nodes), which are candidates to receive a handover when the client moves from one AP to another. This information can also be used for sending clients' authentication profiles to these neighboring APs, so as to minimize the handover time.

With this information, the CC chooses the AP providing the highest AAC for that particular node. In previous sections (c.f. 6.1.5), we argued that AAC was the best load metric for IEEE 802.11, since it captures the effects of all phenomena affecting an AP's load. Each AP computes its own AAC, which is sent to the CC. Then, the CC sends a response to the client (through the cellular interface) with the SSID, channel, encryption key and MAC address of the AP (in addition to the IP addresses and gateway address for the client's WLAN interface configuration).

7.2.3.2.2 Frequency assignment The availability of a GPRS/UMTS interface simplifies the signaling between the CC and the managed APs. The CC receives interference and utilization information that both the APs, and the clients send through the cellular interface. In this way, it is possible to utilize the CC to coordinate the channel allocation management. As previously detailed, the CC receives interference information during the user authentication phase, but also when clients detect a link degradation. Moreover, the mesh nodes send measurement reports in a similar way. After a mesh node is activated and validated against the CC via the GPRS/UMTS interface, the CC sends a message to the AP with channel configuration information. Given a scenario like the one depicted in Fig. 5.3, Fig. 7.5 shows the message sequence sent via UMTS/GPRS between the mesh nodes and the CC, when a new node *2* joins the WMN; nodes *0* and *1* are already operational on channels 1 and 6, respectively. After the first assignment, node *3* (non manageable) is initiated on channel 11:

- *INFO_MSG*: Periodic message reported every T_{info} seconds by the mesh nodes, including information about their identity (MAC$_x$), current channel (Cx), load conditions (U$_x$) and interference information (S$_{xy}$). The packet format is depicted in Fig. 7.6.
- *NEW_CHANNEL_MSG*: unicast message sent by the CC to each of the managed APs with the channel assigned to that node.

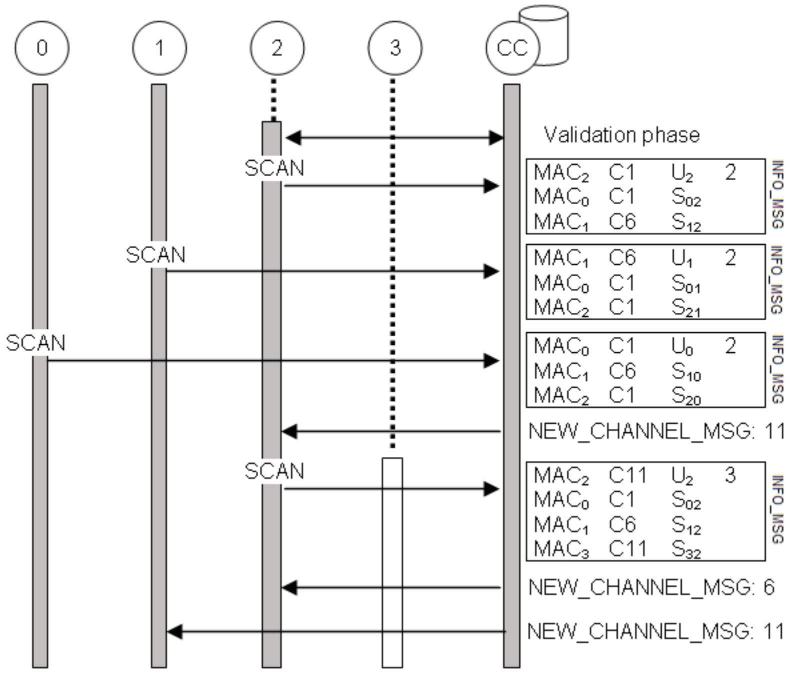

Figure 7.5: Example of signaling exchange for channel assignments in UAMN

All this information allows the CC to build the interference graph, which is required to run the algorithm defined in Section 5.2.2. The CC runs the algorithm periodically (every $T_{valid\_ch\_alloc}$ seconds). If the CC detects that a channel switch could improve the capacity of the network by a preset threshold, it sends the new channel assignment to each AP (only to those APs whose assigned channel must change in the new assignment).

7.3 Distributed management

As discussed earlier, centralized management systems are a good solution for medium-sized networks under a single administrator domain. But they are losing their position due to their poor scalability, and lower levels of robustness and dynamism. Furthermore, more intelligence is continuously being added to all network elements, so they are gaining autonomy and increasing their decision capacity. Next generation networks are distributing core functionalities towards the boundaries of the operator, to the point that 4G anticipates that even client devices will participate [121]. In a scenario where the responsibilities are getting distributed, it is reasonably expected that any problem is faced by promoting the cooperation between the elements that comprise the network. But a distributed and cooperative mechanism always involves some kind of communication between the participant entities; this is the problem we will try to depict in this Section.

Different paradigms exist in the literature for the distributed management. These paradigms include the idea of *mobile code* [86] and its related application known as Management by Delegation (MbD) [201]. The concept of *intelligent agents* was also brought to network management [94], and later applied to IEEE 802.11 WLANs in [125]. However, a concrete definition suitable for the WLAN radio management environment was lacking. For this reason, we developed two new distributed resource management strategies, which are analyzed in this section. These approaches cover the requirements of a channel allocation algorithm (as detailed in Section 5.2) and a distributed load balancing technique based on cell breathing (as detailed in Section 6.2.3). The new communications needs that arise are analyzed. Although we will concentrate on the two mentioned mechanisms, we provide conclusions that remain valid for other distributed resource management strategies.

7.3.1 Requirements for communications

We try to solve the problems of adding signaling capacity to a group of elements which cooperate to obtain a better use of a shared resource. Signaling is a means by which the network elements exchange information to provide control management and performance monitoring (in contrast to user information transfer).

Although the major number of WLAN networks are infrastructure based (use of APs), from an inter-network point of view, its behavior could be compared to that of ad-hoc networks. That is to say, our scenario consists of a varying number of access points that can unpredictably appear and disappear in a region. Moreover, future wireless networks (4G) are expected to have an ad-hoc, dynamic structure with ubiquitous nodes that must have the individual ability to self-configure and promote a collective awareness. Controlling such a network means coping with uncertainty. For this reason, a new approach to the control and coordination of future networks will be needed, replacing centralized with highly decentralized control. A widely accepted approach consists of self-organizing systems found on the interaction of smart but simple nodes providing network-wide coordination. Nevertheless, there is no need to imagine uncertain future scenarios; solving distributed resource management issues in the field of IEEE WLANs is a challenging field of research not only because it is a nearby 4G testbed, but also for actual current applications. Therefore we consider the IEEE 802.11 APs as the entities exchanging signaling information in a distributed manner.

All these arguments led to the implementation and deployment of different distributed resource management strategies in an IEEE 802.11 WLAN ESS. We first considered the necessity of communication among APs when we wanted to implement the distributed algorithm for the frequency assignment problem presented in Section 5.2. For the proper operation of the algorithm, it was necessary to exchange information regarding interference and utilization of every cell, so that all participants know the interference levels between each pair of neighboring APs. Thus, not only is the information exchanged locally among close devices, but we must also provide the means for certain information to travel to the furthest element if needed. In consequence, the elements of the signaling network must be able to forward or route the information in order to reach any AP.

Although its immediate application is limited by the boundaries of a single administrative domain, cooperation among independent domains ruled by different owners is also desirable when they share the same resources, as is the case with the unlicensed 2.4 and 5 GHz frequency bands used in 802.11 WLANs. In these cases, it is clear that cooperation brings a better utilization of the resources than a greedy behavior by the competing domains. But inter-domain cooperative resource management is difficult to establish due to potential mal-practices and other security issues that are out of our scope. We will try to provide means to communicate different domains, although we are aware that the emergent security problems involved are left open.

A load balancing mechanism was also implemented developing the concepts of cell breathing, as detailed in Section 6.2.3. In this case, inter-AP communications are limited to the neighboring set of a congested AP. That is to say, it does not make sense that messages announcing a cell resizing traverse multiple hops. Once the signaling problem is solved for the channel allocation algorithm, it will be easy to add a new message exchange to support the cell breathing technique. The amount of information bytes required for the load balancing is also insignificant compared to the signaling involved in the channel allocation system.

However, the chosen solution should also be suitable for a future use of other resource management mechanisms such as fault recovery, rogue node detection, etc. To sum up, the specific requirements for the signaling protocol are:

- Allows communications between APs in different domains.
- Must follow a distributed architecture.
- Self-configurable, adaptable to changing conditions.
- Should allow all-nodes broadcast and unicast.
- Suitable for different resource management mechanisms.

7.3.2 Feasible solutions

First of all, the lack of a common wired distribution system makes us “think wireless”, i.e. we propose the use of APs with more than one wireless interface: one is devoted to keep working as an AP, and the others can be used to join an ad-hoc network for inter-AP communication. As well as for signaling, inter-AP communication could also be used to build a wireless backbone, as proposed in [299]; in this way, the additional cost introduced by the extra interface could be redeemed. There are solutions that build both

a mesh-based WDS and an infrastructure-based WLAN access network by using single radio devices, as is the case for the products from Firetide and Tropos[4]. However, this approach entails serious performance issues [106].
Furthermore, the absence of wires eases the rapid coverage of communications needs in provisional and temporary locations such as emergencies, sport events, etc.

In the long term, we may be using solutions provided by current standardization efforts. The 802.11s group is creating a standard that will provide mesh features to APs: they will be capable of forwarding data using multi-hop transmissions within an ESS; new IEEE standards, such as 802.11k and 802.11v will enhance management by adopting new measurements and information exchange across layer 2 mechanisms. In a similar way, the current 802.11h amendment offers useful mechanisms for frequency and power management in order to avoid interference with radar systems, but it is not intended to provide communication between distant APs. IEEE 802.11F standard [11] recommends the use of IAPP to provide communication between APs, but since it is mainly focused on the roaming of mobile users and also requires layer 3 mechanisms to route its packets, it will not be considered. Following a different approach, IETF CAPWAP Working Group focuses its efforts in developing a standard for control and provisioning of APs. As detailed in Section 7.2.2, CAPWAP aims to redefine the concept of AP by transferring part of the AP's logic to a centralized Access Controller from which many of these light APs are managed, but this architecture contrasts with the widely accepted future perspectives already discussed. In [133] the simplest solution is implemented: a common wired broadcast domain to which all APs are connected. However, as stated before, if the use of the infrastructure is expensive or there is no common infrastructure at all, this solution is not suitable. The authors in [43] provide the details of an Inter-AP protocol with the aim of implementing a channel assignment mechanism. Their protocol allows this mechanism to be distributed, centralized or semi-centralized (based on clustering). However, they also rely on a common wired backbone.

7.3.2.1 Layer 2 mechanism

Our first proposal is based on a layer 2 mechanism. Each node sends its signaling information in a broadcast frame so it reaches all its one-hop neighbors (i.e. APs within reach of each other). Note that two APs can have overlapped coverage areas and still be out of range of each other; this discussion have been settled in section 5.2.2.1. When a node receives information that needs to be forwarded, it is added at the end of the receiving node's frames and broadcasted again in order to reach the origin node's two-hops neighbors, and so forth.

For the channel allocation application, a new packet is exchanged. The new packet format (see Fig. 7.6) is similar to that presented in [259]: for every neighbor each node detects, they must include a unique identifier (e.g. MAC address), its current channel and

[4] TROPOS Networks http://www.tropos.com

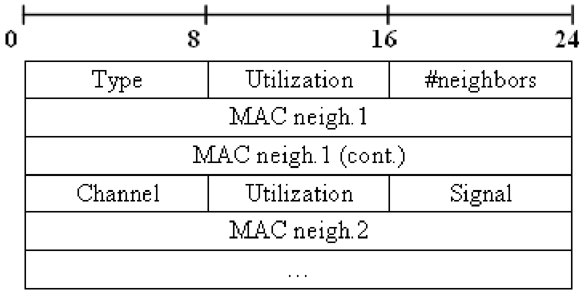

Type	Utilization	#neighbors
MAC neigh.1		
MAC neigh.1 (cont.)		
Channel	Utilization	Signal
MAC neigh.2		
...		

Figure 7.6: New packet format for channel allocation signaling (INFO_MSG)

signal received; we also need cell utilization in order to obtain a more realistic interference assessment (c.f. 5.2). This information is present in the signaling frames as described in the following.
Type field describes the kind of signaling contained in the frame. For a channel allocation frame, the originator node includes its own cell utilization, and the number of interfering neighbors. Next, ID, channel, utilization and signal received from each of its neighbors is included.

Every AP keeps track of all cells and their interference in order to run the channel allocation algorithm. When a new AP is detected, its information is forwarded until it is known by the rest of the APs. Note that this solution does not provide end to end communication beyond one hop neighbors; for that reason it uses flooding at layer 2. In a multi-hop scenario, the flooding of frequency assignment frames could be substituted by the IEEE 802.11s broadcast/multicast services in the near future (see Section 2.5).

The load balancing mechanism requires three different messages: *SOS*, *AAC* and *ACK* (see details in Section 6.2.3) which does not require any further parameter. Since the information is only relevant to 1-hop neighbors there is no need to include any additional field to support forwarding.

7.3.2.2 Layer 3 mechanism

A layer 3 mechanism entails IP level routing, since we must grant communication between IP entities. We can take advantage of current ad-hoc state of the art due to the obvious ad-hoc nature of the problem, so a priori, any well known ad-hoc routing protocol would fit. How quick our scenario changes is the main matter which will settle the choice between the two principal paradigms: *proactive* (table-driven) or *reactive* (on demand). The nodes of our ad-hoc signaling network are WLAN APs; and these networks are assumed to be quite stationary systems. They are also connected to fixed power sources, so energy efficiency is not a key issue to solve. With these premises, the choice becomes proactive. Moreover, since we need to exchange information regarding link status of all neighboring

pairs of nodes (APs) in order to run the channel allocation algorithm, the use of a link state routing protocol may save efforts. All these considerations and a major number of implementations led us to Optimized Link State Routing Protocol (OLSR) [102] rather than other alternatives, such as Cisco's Open Shortest Path First (OSPF) extensions for ad-hoc networks [89] or the Topology Dissemination Based on Reverse-Path Forwarding (TBRPF) protocol [102].
Within the OLSR solution, two different approaches are followed: altering OLSR's own mechanisms to include signal information, or creating a completely new functionality that uses OLSR as a means to route information.

As a link state protocol, OLSR already exchanges information about certain links periodically, but this information only states whether a link exists or not, so we should modify OLSR messages to add signal and utilization data. Furthermore, due to the normal OLSR behavior, not all the existing links are announced through the network. Thus, it will also be necessary to alter the default values of some specific parameters in order to guarantee that all link information arrives at all nodes. In OLSR, each node selects a reduced set of neighbors as "multipoint relays" (MPRs) so that they can reach any 2-hop neighbor using the minimum number of relays. We can add redundancy by configuring more than one MPR for each 2-hop neighbor. This can be done by setting $MPRCoverage > 1$. MPRs are also responsible for forwarding topology information. By default, MPRs only declare link-state information of their MPR selectors (i.e. an MPR only informs about the links with the nodes that have selected it as MPR). The parameter *TCRedundancy* can be modified to force MPRs to announce all their links, not just the links with MPR selectors. This option implies the modification of the standard, and it turns out to be too specific for the channel allocation application (i.e. hardly suitable for other mechanisms).

The second option is based on a new type of OLSR message and the new mechanisms required for dealing with those messages, while the system remains RFC compliant. Every node broadcasts its neighbor information over a standard OLSR enabled network. Even though this kind of message is unknown for a node, it will be forwarded thanks to the default forwarding rule of OLSR. Besides IP and OLSR common headers, the new packet contains the fields depicted in Fig. 7.6 except for Type field, whose information would be carried in the OLSR header. Source MAC and channel should also be included, since it could only be detected by one hop neighbors but it must be known by all nodes.

7.3.2.3 Implementation Details

Two of the proposed solutions were implemented in a practical scenario using commercial Linux-embedded APs (4GSystems AccessCube with two WLAN interfaces; see section 1.2): the layer 2 flooding-based protocol and the layer 3 OLSR functionality. Even though both solutions are implemented following different mechanisms, both are based on the same guiding principles. Figure 7.7 shows the common behavior for both implementations. The first thing to do when the protocol is initiated is to perform a channel scan

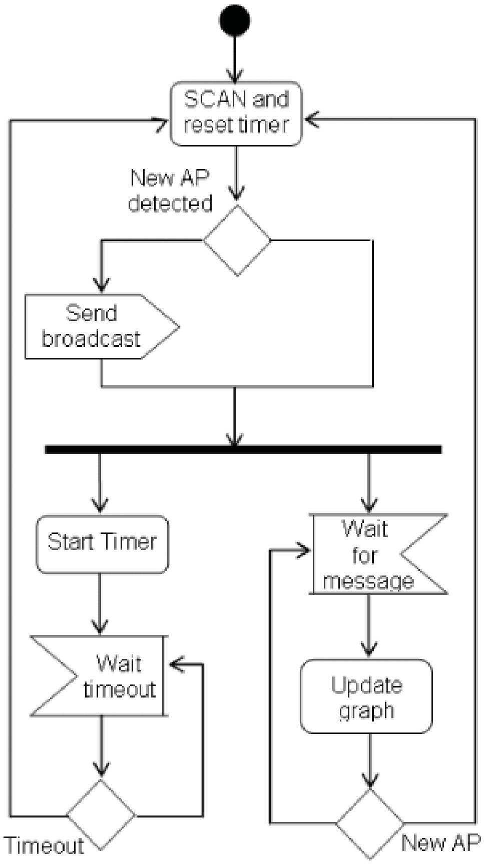

Figure 7.7: Simplified Activity Diagram for graph exchange protocol

in order to obtain information about neighboring cells. Note that two distant APs (out of the decodable range of each other) can belong to interfering cells if there is a client station that is in the transmission range of both APs. In this case, the graph (c.f. section 5.2.3) will not reflect the actual interference, but this issue can be solved by using IEEE 802.11k enabled STAs. Recall that 802.11k allows the APs to request their clients the list of APs in range and other radio measurements that will better describe the actual interference (see Section 7.1.2). If there are two interfering cells whose APs cannot communicate with each other because they are out of each other's transmission range, the needed information will still be shared if there is a path through the signaling network that communicates both APs. If both interfering APs belong to different and independent graphs (i.e. there is no path connecting them), then they will be treated as a non-configurable AP (i.e. it does not run the protocol) in each other's graph.

Recall that frequency spectrum is a scarce resource and a wireless distribution system would add interference to AP radios, but the new signaling channel and the AP interface (devoted to provide connectivity to mobile clients) could use different technologies, and different frequency bands; e.g. IEEE 802.11g for AP cells and 802.11a for inter-AP communications.

Once the ad-hoc signaling network has been detected and configured, a message containing information, as depicted in Fig. 7.6, is broadcasted. Then, periodic scans are re-scheduled so that the information is maintained up-to-date. New messages can be sent asynchronously if the local information has changed significantly (e.g. a new AP is detected, utilization grows above a given threshold, etc.). In parallel, the application remains waiting for other APs' messages. Upon the arrival of a protocol message, the receiving AP updates its interference graph (used by the frequency assignment algorithm), and if new relevant information is learned from that message, a new broadcast message is sent.

In the layer 3 case, a new message containing just local information is sent only when a new AP is discovered in the received message; thus, the new AP gets information about all existing APs' interference. OLSR is responsible for routing these messages to all members.
The absence of a routing protocol in the layer 2 mechanism requires a different behavior to ensure the delivery of information beyond one-hop neighbors. Either when new information is learned from a protocol message, or when known information has significantly changed, besides local information, the next message sent must contain all changes detected in the graph.

The load balancing mechanism for both implementations follows the diagram shown in Fig. 6.20.

7.3.2.3.1 Layer 2 mechanism The layer 2 protocol was implemented in *C* using the Linux networking socket layer's low level packet user interface. Type *SOCK_RAW* sockets of *AF_PACKET* family are used to receive or send raw packets at the device driver level. These sockets allow the user to implement protocol modules in user space on top of the physical layer. All incoming packets of that protocol type will be passed to the socket before they are passed to the OS kernel. Therefore, the kernel task of generating and interpreting physical and link layer headers must now be carried out by the user application. Packets are then queued unmodified to the network driver of the interface defined by the destination address.

Figure 7.8 represents the message sequence that the layer 2 mechanism will follow in the scenario depicted in Figure 7.10.2. In this example, cell B is ruled by an AP which is not implementing the protocol and is tuned to channel 11. When C is initialized, it first performs a frequency scan after which it decides to select channel 1, since no other neighboring cell has been detected. The next AP to start up is A; after the scan, A broadcasts information about its cell utilization (U_A) and the signal level received from cells B and C (S_{BA} and S_{CA}). This message is responded to by C, thus completing all the required information regarding $A \leftrightarrow C$ link. A selects channel 6 (1 and 11 are in use

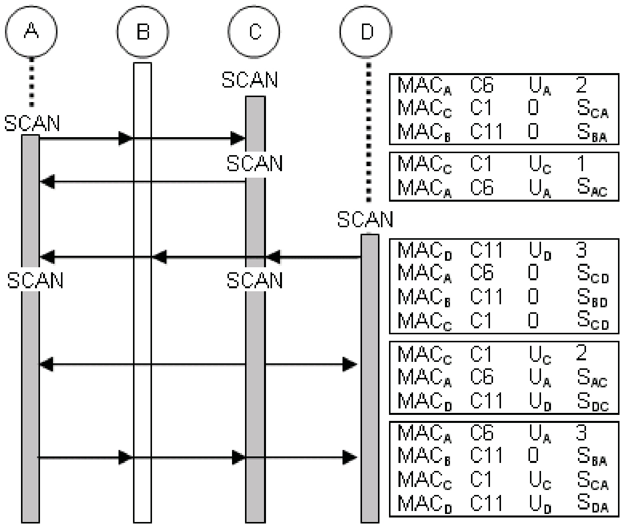

Figure 7.8: Example of signaling message sequence for L2 protocol

by its neighbors). Since *C* does not know of any other node that is not included in *A*'s message, there is no need to forward that information. *D* is initiated; after a scan, *D* realizes that all non-overlapping channels are in use, selects channel 11 and broadcasts information about its three links; the message is responded to by *A* and *C*. The absence of messages from *B* makes the rest of the nodes tie *B* with channel 11. At this point, knowing the interference graph, the channel allocation algorithm assigns channel 6 to *A*, channel 11 to *B* (fixed), channel 11 to *C* and channel 1 to *D*, avoiding all interference among the cells of the graph. If a new AP appears creating one link *D*↔*E*, *D* would be responsible for forwarding information about nodes that are not linked with *E*.

7.3.2.3.2 Layer 3 mechanism Unik-OLSR[5] is one of the most popular open source OLSR implementations. The code is structured in such a way that allows the programming of new functionalities by adding plugins. These plugins are dynamically linked libraries (DLL), known as *.so* in UNIX systems. The basis of this implementation is the OLSR service daemon, *olsrd*. The modular design of the various *olsrd* entities provides an easy way to add special functionalities into most aspects of the protocol. This main feature really comes to the fore when working with plugins (see Fig. 7.9). New functionalities are added by registering custom functions with the socket parser or the packet parser. The scheduler module is used to periodically broadcast the required information. More details of Unik-OLSR plugin interface can be found in [275].

The implementation of this inter-AP signaling using OLSR is detailed in [116]. The plugin designed for this implementation has been written in C, though it could be done in

[5] http://www.olsr.org

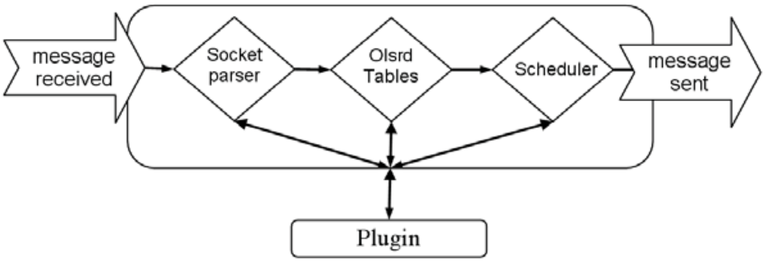

Figure 7.9: OLSR plugin interaction interfaces

any language that can be compiled as a dynamic library. OLSR plugins are used to achieve two main goals: creating new messages to add new functionalities, or changing the normal OLSR behavior. Our plugin has been developed to create a new type of OLSR packet that is broadcasted across the network, using the *olsrd*. In this case a new functionality is needed, and a new message must be created, following the OLSR standardized structure. OLSR messages are identified by a number in the header: numbers 0 to 127 are restricted, but 128 to 255 are free and available for additional new features. This new type of message will transport the information described in Section 7.3.2.2.

For the channel assignment mechanism, the AP creates a message with information about the configuration and utilization of its own cell, along with data including interference power received from its neighboring nodes. These messages will be forwarded by any node using standard OLSR mechanisms, so developers do not need to bother about implementing this functionality. Besides the periodic broadcast of local information, the OLSR daemon with our plugin follows a simple workflow upon reception of this new type of message. AP first checks whether it contains information about a new AP or not. If it does, a neighbor scan is performed in order to know if the new cell is an interferer, in that case, a new broadcast is triggered so all APs know the new interference link. If the new node is not an interferer, a unicast message is sent to that AP so it can build the complete graph.

7.3.3 Performance tests

With the purpose of evaluating the performance of the formerly proposed solutions (i.e. layer 2 and layer 3 distributed mechanisms), we measured the amount of signaling information as the number of participating APs is increased. Results were obtained through simulation and, due to their simplicity, two of the proposed solutions were fully developed and tested. Besides practical results, the implementations helped to show more precisely the hardware, software and knowledge requirements for the deployment of different solutions; those details were discussed in Section 7.3.2.3. From now on we will concentrate on interference information exchange to improve channel allocations, since as it detailed

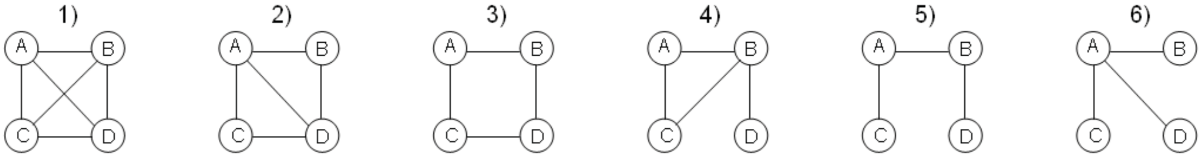

Figure 7.10: All possible topologies built from a connected graph with 4 nodes

through previous sections, it involves a greater complexity than the cell breathing mechanism, due to the need of multihop forwarding and a variable message size.

A first set of simulations was intended to study the feasibility of the OLSR-based approach (modification of standard parameters and standard messages). Several topologies were created and simulated using OOLSR simulation tool developed by Hipercom at INRIA[6]. At first, it was logical to think that, depending on the topology, not all interference links would be announced through the network even though we tune *TCRedundancy* and *MPRCoverage*. For example, let the topology under test be like that in Fig. 7.10.4: all nodes chose B as their MPR. Theoretically, with a *TCRedundancy* set to 2, B broadcasts topology control messages announcing its entire link set ($A \leftrightarrow B$; $B \leftrightarrow C$ and $B \leftrightarrow D$). Thus, the link $A \leftrightarrow C$ will not be known by D, causing a malfunction of the channel allocation algorithm. But in the practice, those simulations (and practical testbeds too) showed that all interference links were actually announced throughout the network, due to the fact that all nodes, independently of their MPR condition, sent TC messages announcing all their links. We have to note that this fact was interpreted as an “inconsistency” of OOLSR implementation, but it was later verified using Unik-OLSR implementation. Anyway, the cost of signaling is increased severely since the efficiency gained with the use of MPRs is lost due to the excess of redundancy.

In order to measure the amount of signaling information, a simulator developed in Java was implemented. This tool simulates the exchange of signaling messages following the Layer 2 broadcast mechanism, the OLSR plugin and a Layer 3 broadcast over a common wired distribution system, as introduced in [133]. The application randomly places the APs over a flat grid so that an undirected and connected interference graph is achieved. An undirected graph is considered connected when there is a path between every pair of nodes. If the interference graph were not connected, it could be decomposed in independent subgraphs which could be solved separately. Then, following a random sequence, each node of the graph is activated, starting the signaling protocol. The amount of Bytes sent and received by all nodes is measured at application layer, taking into account the payload carried by MAC frames (i.e. IP, and UDP headers are counted in L3 broadcast; and IP, UDP and OLSR headers in OLSR plugin approach). Also note that inherent OLSR signaling (i.e. Hello and TC messages), which is not an explicit component of our protocol, is not considered. Its addition would increase the bytes of signaling information, but the

[6] http://hipercom.inria.fr

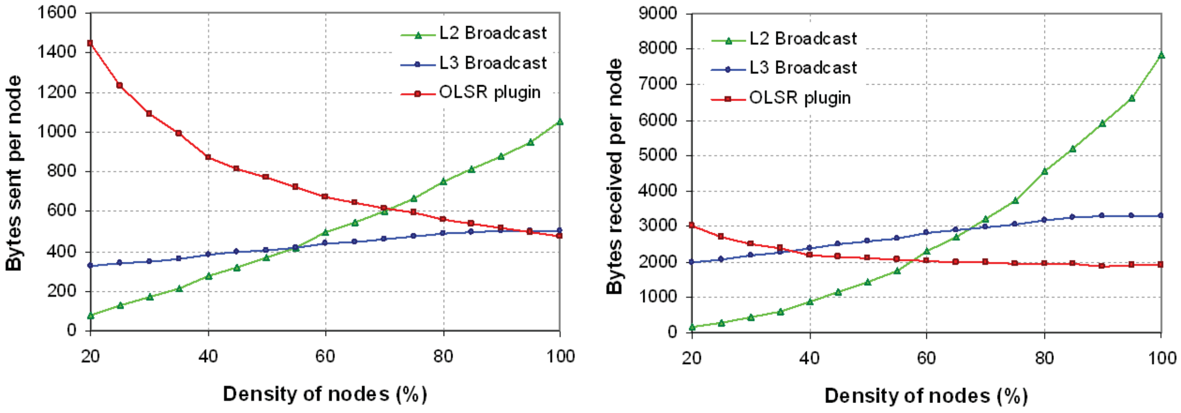

(a) Bytes sent by each AP

(b) Bytes received at each AP

Figure 7.11: Signaling Bytes for 10 APs as a function of density

general trends observed in the figures would remain without variation.

Simulations with 10 nodes and a varying density were run. In Figure 7.3.3, different densities are represented in %, which is interpreted as the portion of edges the generated graph has, compared to a complete graph with the same number of nodes. In a complete graph, there is an edge connecting every pair of nodes (1-hop path between every pair of nodes). The number of edges in a complete graph with n nodes can be computed as $n(n-1)/2$. A density of 20% with 10 nodes means that there are 9 edges in the interference graph.
From the figures it is clear that the L2 mechanism is the most efficient solution with lower densities, but it is the worst under higher densities. In contrast, the OLSR mechanism improves if density is increased due to the reduction in the number of hops needed to communicate distant nodes. Even though there is always a single hop through the common wired distribution system in the L3 broadcast, it is outperformed by the OLSR mechanism, using a multi-hop topology, thanks to the use of unicast for certain messages, as explained in Section 7.3.2.3. Recall that in the wired distribution system all messages are sent to the broadcast address.

Fixing a density of 65% (similar performance for 10 nodes in all three solutions) and varying the number of participant APs, new simulations were executed (see Figure 7.3.3). It is again evident that L2 broadcast mechanism is the best solution for small networks, and at the same time, the lack of scalability is noticeable. Therefore, some modifications of the algorithm should be studied for its implementation in large scale scenarios. It was also found that the total number of broadcast messages sent in the L3 broadcast mechanism and the total unicast messages sent in the OLSR mechanism grow as $O(n^2)$, while the number of broadcast sent through OLSR increases linearly with the number of edges, hence producing a lesser impact.

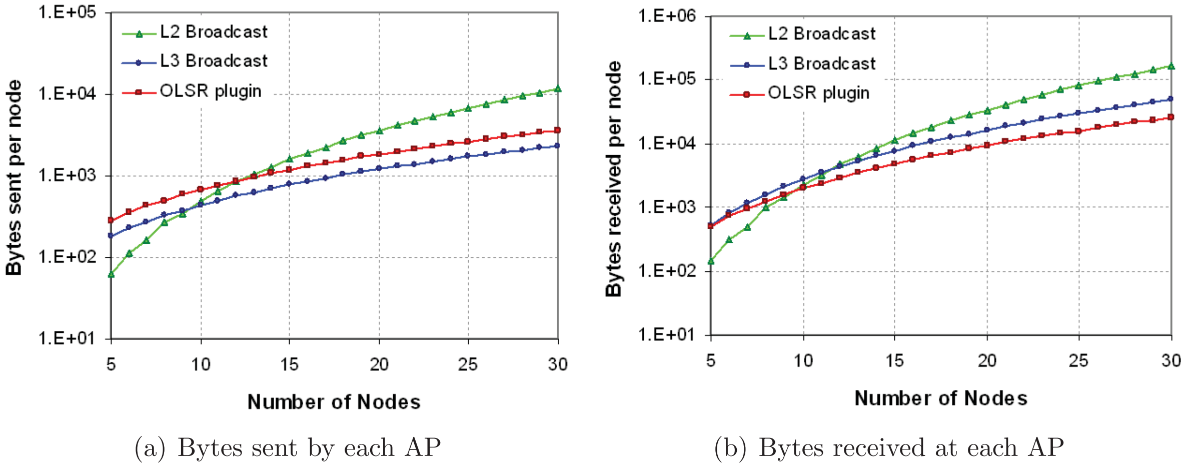

(a) Bytes sent by each AP

(b) Bytes received at each AP

Figure 7.12: Signaling Bytes for fixed density (65%) as a function of the number of APs

Note that these results represent the signaling cost produced during the constitution of the mesh network. Once the network is established, the signaling is reduced. Signalling messages are triggered by events such as mobility, the presence of new APs and their removal. According to [45], such events rarely occur and therefore small differences are observed between the different solutions in the long run. In the event of a topology change, the wire-based Layer 3 broadcast requires fewer bytes to spread the news than the other solutions, followed by the OLSR plugin.

Finally, a testbed, consisting of 4 overlapping cells was built. Different topologies were tested (see Fig. 7.10) in order to compare convergence times of the three solutions. The time was measured after the last AP was activated, and until all of them were able to build the entire interference graph. As expected, the L3 broadcast over a common wired distribution system outperforms the other two. In the other two solutions, the information traverses multiple hops, adding an extra delay. Moreover, the OLSR plugin needs a previous convergence of the routing protocol so that all APs are reachable. The average times for different scenarios are: 4.5s for L3 broadcast, 5.8s for L2 broadcast and 8.2s for OLSR plugin. Measured convergence times are fast enough if we consider that interference does not vary significantly for long periods, and hence channel assignments remain valid for intervals of tens of seconds or even minutes. As argued in Section 5.2, the exchange of information was crucial to dynamically tune to the best assignment, obtaining throughput improvements of 15-20%.

7.4 Conclusions

In order to bring the radio resource management techniques described in previous chapters into practice, a study on the different management architectures was needed.

First of all, we could take advantage of new functionalities that are already present, or planned, on newer IEEE 802.11 devices. For this reason, this chapter provides a study on the new IEEE 802.11 standards.

Then, a centralized and a distributed approach are discussed. In the first case, APs are managed from a central point under the control of the network operator. Signaling is exchanged by using a UMTS/GPRS cellular interface. Note that this approach uses the cellular interface as a highly available Internet connection that allows the communication with the central unit; i.e. the system is independent of the cellular network operator and in consequence it can be implemented by any kind of provider. Nevertheless, the project also considers the possibility of taking advantage of the operator's AAA mechanisms through additional agreements.

Centralized management mechanisms are well known and do not involve extra complexity. These paradigms are therefore easy to implement. However, given the great diversity of WLAN networks that may coexist, it may be difficult to coordinate all nodes in order to optimize the utilization of radio resources. This makes centralized management solutions difficult to put into practice. A distributed system of some sort would therefore be appropriate, but this would require additional intelligence for the participating devices. This requirement is supported by current technology trends, which are expected to keep producing devices with improved features, i.e. new APs will have faster CPUs, greater storage/memory capacity, etc.

Distributed management also raises the requirement of signaling among the participant nodes. A signaling protocol could assist many resource management mechanisms to improve the performance of a network. We have introduced a key issue involved in providing means for a distributed resource management algorithm for WLANs that cannot be connected through a common wired distribution system. Signaling among APs is mainly achieved by adding an extra wireless interface for ad-hoc communications. We have also detailed three possible solutions whose performance is evaluated by means of simulation, implementation on commercial APs and practical results. The use of OLSR plugins eases the development of a signaling protocol and has been verified as the most efficient solution, even though it has the slowest convergence time. A layer 2 approach performs well in the presence of a small number of cells, but it should be revised to fit into large scale scenarios.

Chapter 8

Final discussion

The present chapter concludes this dissertation and is therefore intended to summarize not only the course of this project, but also to foresee future steps. More precisely, this chapter presents the concluding remarks of the research developed during this thesis, including our main contributions. Furthermore, some future lines of research are anticipated by means of a brief state of the art review.

8.1 Conclusions

From the introduction and until the end of Chapter 7, one thing is made clear: IEEE 802.11 WLANs were not designed to bear their own success. This means that the processes and mechanisms that the IEEE 802.11 standards defined were not conceived for the highly competitive and overloaded environment in which they are usually deployed. This is further compounded by the use of unlicensed frequency bands (ISM). In consequence, IEEE 802.11 WLANs do not perform efficiently in dense scenarios, mainly due to contention and interference. There are two possible ways to cope with these issues: by re-defining the PHY and MAC layers, and by implementing radio resource management mechanisms. Refurbishing the standards was out of our scope, so we chose to define new radio resource management strategies in order to minimize the impact of interference and the effects of a high level of contention. Throughout this thesis we have explored several ways of reducing such harmful effects in the particular field of IEEE 802.11 WLANs. In order to reduce interference and contention, we propose a new channel assignment scheme. We also demonstrate that contention can be further reduced by means of an efficient load balancing strategy, which can be combined with power control. The proposal of those approaches required a previous study of all the issues involved; such analyses are also part of the contributions of this thesis.

Frequency management requires the understanding of the interference. The first conclusion drawn from an early analysis was the conviction that a new approach for modeling interference in WLANs was needed, mainly due to the use of a CSMA/CA access. In this regard, we decided to study adjacent channel interference and co-channel interference separately. The first is produced by transmissions in adjacent or partially overlapping channels. These transmissions are not perfectly filtered by the receivers, and are thus added to the desired signal. In this way, the SINR is degraded, and as a consequence, the BER is increased. In other words, bit and frame errors are produced in reception at a higher rate. In Chapter 3 we model these phenomena by means of a novel approach that takes into account not only the channel distance between the interference source and the receiver, but also the interferer's activity (i.e. its utilization), thus capturing the actual overlapping degree between the desired and the interfering signals. Adjacent channel interference, together with noise are finally translated into a PER value.

On the other hand, co-channel interference is produced by simultaneous transmissions in the same frequency channel, or either by transmissions in partially overlapping channels (if received with enough energy). Co-channel interference is then determined by the number and activity of stations sharing the same channel, i.e. the level of contention. The presence of co-channel interference prevents an IEEE 802.11 station from transmitting, due to the "listen before talk" strategy that the CSMA/CA type of access defines. Moreover, co-channel transmitters contribute to increasing the collision probability, and may be the source of the hidden node problem (c.f. Section 3.3.5).

Although it is done in different ways, the two types of interference affect the network performance by reducing its capacity. Both co-channel and adjacent channel interference effects are jointly studied in Chapter 4, where we provide a model that allows the estimation of the overall capacity in an IEEE 802.11 ESS. These estimations are useful to quantify the effects of the interference, and can thus be used as the target metric to be optimized by a frequency channel allocation mechanism. As a matter of fact, our approach is based on these capacity estimations, as described in Chapter 5. In the scheme we propose, an interference graph is built, which includes information on the utilization and the interference received by all the 802.11 AP. Then, a channel assignment is decided by following an adaptation of the well-known DSATUR graph coloring algorithm. This adaptation consists in adding weights to the edges of the graph, depending on the degree of interference produced by the different 802.11 cells. This approach does not always provide the optimal channel allocation, but on the other hand, it is fast and allows for a timely adaptation to changes in the environment. Besides, the simplicity of graph coloring eases its implementation in a distributed architecture, where the APs, which have limited memory and process capacity, can run the algorithm without sacrificing excessive resources.

Thanks to the fact that we are able to model adjacent channel interference, we can introduce another remarkable novelty in our solution: the use of all the available spectrum. Most of the literature was based on the use of non-overlapping channels; recall that in the case of the 2.4 GHz ISM band, this is limited to just three channels. We demonstrate that in dense scenarios, the use of all the available channels outperforms the traditional approaches.

As explained, contention is reduced by minimizing the number of neighboring cells that use the same channel. Nevertheless, if an IEEE 802.11 cell is serving a large number of active stations, the levels of contention will still be high. Several surveys stated that Wi-Fi users tend to concentrate in *hot-spots*, and given that client stations exchange their traffic with the "closest" AP, this produces an uneven load distribution: the APs that are located closer to the herd of Wi-Fi users will have to support a heavy load, while distant APs will be under-utilized, even though they are capable of serving some of the users in the hot-spot. Logically, the available throughput per station is inversely proportional to the number of contenders, since they have to compete for a shared channel; and besides, the collision probability is increased. In Chapter 6 we argue that this situation can be mitigated by redistributing the load among neighboring APs in a fairer way. However, it was first necessary to define the concept of load. A study on different parameters led us to the definition of the *available admission capacity* (AAC). This measurement is proved to capture effectively all the effects that influence the actual load of a cell: signal quality, number of stations, multi-rate transmissions, collisions, packet losses and offered traffic. Based on this definition of load, two different load balancing strategies are proposed in this thesis. The first one consists in changing the criterion followed by client stations when deciding the best candidate AP: instead of considering only signal strength measurements, clients choose the AP providing a higher AAC. The second approach is based on the cell breathing concept. This comprises APs with the ability to dynamically changing their cell radius according to their neighbors' load, and of course, their own load. An overloaded AP will reduce its coverage (i.e. will reduce transmission power), while the neighboring (and under-utilized) APs will increase their power in order to enlarge their cell, hence gaining more clients. Both approaches have shown to improve the network performance in terms of global capacity and fairness, and both solutions are based on the standard client-driven association management; that is, roaming decisions remain on the client side, as stated by the IEEE 802.11 standard.

The goodness of our approaches and the benefits of our proposals have been proved by means of mathematical analysis and simulations. Moreover, in the pragmatic spirit that characterizes our research group, we also set out to prove that our ideas were not only practical, but also practicable. To this end, those ideas have been brought to life by applying certain modifications to off-the-shelf equipment (see Section 1.2). We should point out that not all the solutions were practicable with the currently available technology, but as reviewed in Section 7.1, recently approved standards and other ongoing studies will solve those issues in the short term.

With the implementations, new problems arise. In order to integrate radio resource mechanisms into a management system, it is required that we define an embracing architecture. In this thesis we explore two paths: distributed and centralized. We have seen that both paradigms have pros and cons in terms of simplicity, scalability, reliability, etc. For example, centralized management is not suitable for chaotic scenarios (coexistence of uncontrolled WLANs from different owners), but on the other hand, it is easier to implement. In Chapter 7, these two approaches are presented. The centralized architecture, namely UAMN, fulfills the requirements of an operator-centric solution. The distributed architecture supports the presence of non-participating devices and allows interaction between different domains. However, the signaling used to exchange the required information (e.g. to build the interference graph) is more complex, and the algorithms need to be simplified, since they run on low-featured devices, while a centralized management can use high performance computers.

All in all, we have the arguments to claim that the mechanisms proposed in this thesis, which are the result of a deep study of the phenomena involved, achieve their purpose: the network performance, and consequently the user experience, is improved by means of an intelligent and automatic radio resource management. The main contributions of this work, which have been cited by other authors, have been published in notable scientific journals and in well-known conferences and symposiums. Beyond these academic achievements, the development of this thesis has been more than just an enriching experience: publications entailed a direct feedback from other professionals, who helped to improve our methods and from whom we discovered different perspectives. Furthermore, conference speeches provided the opportunity to defend our approaches before an expert audience in international forums. Moreover, struggling with IEEE standards enabled us to acquire priceless knowledge about wireless networks (and has made me immune to otherwise painful technical documentation). The experience gained in the process of putting our ideas into practice and facing real-life problems can also be considered a positive attribute of our work.

8.2 Future Work

The contents of this final section have two objectives: first, we would like to detect the weaknesses in our work and to discuss possible ways of improving our approaches; secondly, we will try to describe the evolution of this research, taking into account the state of the art.

At present, we are still studying the effects of interference on IEEE 802.11 networks. In this case, we are investigating the effects of a broadband jammer and its influence on an 802.11 transmission. Based on this study, we are developing a new mechanism to detect the presence of a jammer from an IEEE 802.11 AP. We will subsequently try to minimize the impact of such attack by means of the load balancing techniques proposed in this thesis.

Also directly related to the contents of this dissertation, the joint study of channel and power control is probably the most evident unresolved matter in this work. Therefore, it would be the natural next step forward. Even though the effect of transmitted power has obviously been considered within our interference model, and has thus been taken into account in our channel allocation mechanism, a unified strategy that considers simultaneous frequency and transmission power settings has not been studied. However, the benefits of such a joint mechanisms can be anticipated from the evaluation of our cell breathing approach, where an upper bound for this mechanism's improvements is established, given that the evaluation is based on the assumption of an optimal channel assignment (i.e. there is no inter-cell interference).

On a slightly different topic, new IEEE 802.11 activities draw our attention. Even though at the time of writing it is still under development, the IEEE 802.11n [39] provides several new challenges to RRM researchers. IEEE 802.11n promises PHY rates up to 600 Mbps, through an overwhelming array of combinations of features and options, yielding 576 different PHY data rates (recall that IEEE 802.11b had only 4 modes!). These improvements are mainly due to the use of MIMO techniques. Among other benefits, MIMO is used to exploit spatial multiplexing (SMX), that is, multiple streams (independent and separately encoded data signals) are transmitted simultaneously from each of the multiple transmit antennas, using the same frequency channel. Another MIMO-enabled feature is beamforming, which improves the received signal quality at the decoding stage. These advances completely change the way interference is understood. In addition, 802.11n offers *channel bonding*, with which a transmitter can simultaneously use two separate non-overlapping channels to transmit data. This allows direct doubling of the PHY rate from a traditional 802.11's 20 MHz channel, but it doubles the spectrum required, and at the same time, the need of RRM.

Tracking the list of new incorporations into the IEEE 802.11 family, another interesting issue arises with the recent release of the IEEE 802.11y amendment [35]. Recall from Section 2.5 that IEEE 802.11y enables high power transmissions on licensed bands. This novelty alone would justify our interest, but the IEEE 802.11y also introduces changes to the carrier sensing and energy detection functions of 802.11, and extends channel management mechanisms previously defined by IEEE 802.11h, and later used by 11k and 11n.

Nevertheless, the future of radio resource management in WLANs is undoubtedly tied to another concept that will be present in next generation communications: *Cognitive Radio (CR)*. This concept, introduced by Mitola in [227], involves adding new intelligence to the lowest OSI layers in order to provide better adaptations to the environment. The new transmitters must be aware of the evolution of the channel and take decisions based on the observation of the medium. According to this information, some actions are taken to meet certain pre-defined goals, either individually or in cooperation with other devices. These ideas require a technology that is not fully available as yet. However, they are being applied gradually, starting with pre-programmed systems, generally based on Software Defined Radio (SDR) [226].

An SDR transceiver performs most of the signal processing through software: modulation, coding, filtering, amplification, etc. Note that, for the sake of efficiency, these functions are traditionally carried out on hardware. The software approach not only allows an easy adaptation of one or two parameters, but also enables a complete refurbishing of the OSI layers 1 and 2. This flexibility makes the SDR a very attractive platform for the implementation of CR.

The first practical application of the CR focuses on Dynamic Spectrum Allocation (DSA) techniques, which are intended to take advantage of under-utilized portions of the spectrum [156]. A survey on CR techniques applied to DSA can be found in [47].

Lately, CR has been brought to sensor and mesh networks (e.g. see [95]), though in a simplified version, due to the limited features of those devices. Moreover, with the recent publication of the IEEE 802.15.4a standard [25], Ultra Wide Band (UWB) becomes a member of the WPAN family. UWB uses large portions of the spectrum (500 MHz) with very low power density. The combination of these technologies (CR, UWB, mesh) provides new research challenges:

- Strategies for the discovery of the wireless channel: awareness of the environment, individual or through cooperation.
- Solving conflicts: the goals of different users can be opposed. The solution to these conflicts is usually modeled by applying game theory [148].
- Waveform adaptation: different access networks can coexist by changing their transmitted waveforms in real-time [65, 148]

Moreover, cooperative diversity and coding techniques could be applied. For example in [216], the CR is also aware of other phenomena, such as propagation (multipath effects, noise, etc.). In this way, frequency bands especially affected by fading could be detected and avoided, making transmissions more energy efficient.

From a more general perspective, the use of CR entails a cross-layer scheme [273]. This vision is also provided in other initiatives; for example, the ARAGORN project [213] shares the idea that CR concepts must not be relegated solely to the lowest OSI layers. In this way, an intelligent radio resource management system is aware of the environment not only at the physical layer (interference, active nodes, occupied channels, etc.), but also by observing other parameters, such as the traffic load, latency, packet loss, etc. Therefore, an optimal adaptation to the environment covers modulation, coding, and channel access issues, and besides, it influences the route selection (layer 3), transmission window (layer 4), or even the selection of an adequate video codec, as we propose in [131].

The cross-layer concept (CL) violates the traditional relationship among layers within a network architecture, since CL enables direct communications between non-adjacent layers [271]. However, this violation is required in order to apply the CR concept at a global level. For example, in [245] a multi-objective optimization is achieved (setting parameters of different layers) by means of a cross-layer design and the learning capacity of the system. In [214], a complete review of CL architectures applied to the optimization of wireless networks is provided.

Adding cooperation capabilities to those network nodes makes the C-Cube concept (C3) [53] to bloom: CL + CR + Cooperation. C3 techniques have an evident application in the RRM field, and a wide spread of approaches can be found in the literature (e.g. see some general remarks on [67, 214]). However, its application on wireless mesh networks presents new challenges. For example, the development of new routing metrics based on layer 1 and 2 statistics can improve network throughput and load balancing. In this way, the best path is not necessarily the one containing less hops, but could be the one using the fastest modulations or suffering less interference, etc. (e.g. see [214, 254]).

Bibliography

[1] OPNET Modeler.
http://www.opnet.com/solutions/network_rd/modeler.html.

[2] OPNET Modeler Wireless Suite.
http://www.opnet.com/solutions/network_rd/modeler_wireless.html.

[3] IEEE Standard for Information Technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. ANSI/IEEE Std ISO/IEC 8802-11:1999(E), IEEE, August 1999.

[4] Supplement to Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 1: High-speed Physical Layer in the 5 GHz band. ANSI/IEEE Std ISO/IEC 8802-11:1999/Amd 1:2000(E), IEEE, 1999.

[5] Supplement to Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Higher speed Physical Layer (PHY) extension in the 2.4 GHz band. ANSI/IEEE Std 802.11b-1999, IEEE, September 1999.

[6] GloMoSim: Global Mobile Information Systems Simulation Library, v2.0.
http://pcl.cs.ucla.edu/projects/glomosim/, December 2000.

[7] HIPERLAN Type 2; System Overview. Tech. Report TR 101 683 V1.1.1, ETSI BRAN, February 2000.

[8] Supplement to Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 2: Higher-speed Physical Layer (PHY) extension in the 2.4 GHz band-Corrigendum1. ANSI/IEEE Std 802.11b-1999/Cor 1-2001, IEEE, November 2001.

[9] HFA3861B; Direct Sequence Spread Spectrum Baseband Processor. Data sheet, Intersil, February 2002.

[10] PRISM Driver Programmer's Manual. (For distribution under NDA only) version 2.30, Intersil, June 2002.

[11] IEEE Trial-Use Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution Systems Supporting IEEE 802.11 Operation. ANSI/IEEE Std 802.11f-2003, IEEE, July 2003.

[12] Part 15.2: Coexistence of Wireless Personal Area Networks with Other Wireless Devices Operating in Unlicensed Frequency Bands. IEEE Std 802.15.2-2003), IEEE, August 2003.

[13] Supplement to Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 4: Further Higher-Speed Physical Layer Extension in the 2.4 GHz Band. ANSI/IEEE Std 802.11g-2003, IEEE, June 2003.

[14] Supplement to Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 5: Spectrum and Transmit Power Management Extensions in the 5 GHz band in Europe. ANSI/IEEE Std 802.11h-2003, IEEE, October 2003.

[15] The Effects of Adjacent Channel Rejection and Adjacent Channel Interference on 802.11 WLAN Performance. White paper, Texas Instruments, November 2003.

[16] on the harmonised use of the 5 GHz frequency bands for the implementation of Wireless Access Systems including Radio Local Area Networks (WAS/RLANs). ECC Decision ECC/DEC/(04)08, Electronic Communications Commitee, November 2004.

[17] Orinoco ap-600 access point. Data sheet, Proxim Corp., 2004.

[18] Supplement to Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 6: Medium Access Control (MAC) Security Enhancements. ANSI/IEEE Std 802.11i-2004, IEEE, July 2004.

[19] Cisco aironet access points features. Url:
http://www.cisco.com/asiapac/powernow/wlan/product_feat.shtml, Cisco Systems Inc., 2005.

[20] Handheld Spectrum Analyzer R&S FSH. Data sheet, Rohde&Schwarz, June 2005.

[21] ns-2 Network Simulator - Contributed Module INRIA-Planète Group, Version 14.2, September 2005.

[22] Supplement to Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements. ANSI/IEEE Std 802.11e-2005, IEEE, November 2005.

[23] Electromagnetic compatibility and Radio spectrum Matters (ERM); Wideband transmission systems; Data transmission equipment operating in the 2,4 GHz ISM band and using wide band modulation techniques; Harmonized EN covering essential requirements under article 3.2 of the R&TTE Directive. Harmonized European Standard (Telecommunications series) EN 300 328 v1.7.1, ETSI ERM TG11, October 2006.

[24] Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. ANSI/IEEE Std IEEE Std. 802.11-2007, IEEE, June 2007.

[25] Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (WPANs). IEEE Std 802.15.4a-2007, IEEE, August 2007.

[26] Receive Sensitivity: A Practical Explanation. Technology brief, Tropos Networks, July 2007.

[27] 5 GHz high performance RLAN; Harmonized EN covering the essential requirements of article 3.2 of the R&TTE Directive . Harmonized European Standard (Telecommunications series) EN 301 893 v1.5.1, ETSI BRAN, April 2008.

[28] AirWave Management Platform. Data Sheet DS_AWMP_US_081117, Aruba Networks, Inc., November 2008.

[29] Cisco Wireless Control System (WCS). Data Sheet C78-341404-12, Cisco Systems, Inc., December 2008.

[30] Intel® Wireless WiFi Link drivers for Linux. http://intellinuxwireless.org/, 2008.

[31] ITU-R Radio Regulations. Technical report, International Telecommunication Union Radiocommunication Sector (ITU-R), Geneva, Switzerland, September 2008.

[32] The ns-3 network simulator.
http://www.nsnam.org/, November 2008.

[33] OMNeT++ Discrete Event Simulation System, V4.0b8.
http://www.omnetpp.org/, November 2008.

[34] QualNet Developer, v4.5.
http://www.qualnet.com, 2008.

[35] Specific equirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 3: 3650–3700 MHz Operation in USA. ANSI/IEEE Std 802.11y-2008, IEEE, November 2008.

[36] Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 1: Radio Resource Measurement of Wireless LANs. ANSI/IEEE Std 802.11k-2008, IEEE, December 2008.

[37] Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 9: Wireless Network Management. IEEE Draft Std P802.11v/D3.01,, IEEE, December 2008.

[38] AT&T Sees Surge in Wi-Fi Connections. http://www.att.com/gen/press-room?pid=4800&cdvn=news&newsarticleid=26765, April 2009.

[39] Specific equirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 5: Enhancements for Higher Throughput. IEEE Draft Std P802.11n/D11.0, IEEE, June 2009.

[40] Wireless Management Suite. Data sheet, Enterasys Networks, Inc., March 2009.

[41] K. I. Aardal, S. P. M. van Hoesel, A. M. C. A. Koster, C. Mannino, and A. Sassano. Models and Solution Techniques for the Frequency Assignment Problem. *Annals of Operations Research*, 153(1):79–129, 2007.

[42] K.I. Aardal, S. P. M. van Hoesel, A. M. C. A. Koster, C. Mannino, and A. Sassano. Models and Solution Techniques for the Frequency Assignment Problem. *4OR*, 1(4):261–317, 2003.

[43] Murad Abusubaih, James Gross, and Adam Wolisz. An Inter-Access Point Coordination Protocol for Dynamic Channel Selection in IEEE802.11 Wireless LANs. In *1st IEEE Workshop on Autonomic Communications and Network Management, ACNM'07*, pages 17–24, May 2007.

[44] N. Ahmed and S. Keshav. SMARTA: a self-managing architecture for thin access points. In *Proceedings of the 2006 ACM CoNEXT conference, CoNEXT'06*, pages 1–12, New York, NY, USA, 2006. ACM.

[45] Aditya Akella, Glenn Judd, Srinivasan Seshan, and Peter Steenkiste. Self-management in chaotic wireless deployments. *Wireless Networks*, 13(6):737–755, 2007.

[46] R. Akl and A. Arepally. Dynamic Channel Assignment in IEEE 802.11 Networks. In *IEEE International Conference on Portable Information Devices, PORTABLE'07*, May 2007.

[47] AIF. Akyildiz, WY. Lee, MC. Vuran, and S. Mohanty. NeXt generation/dynamic spectrum access/cognitive radio wireless networks: a survey. *Computer Networks*, 50(13):2127–2159, 2006.

[48] H. Al-Rizzo, M. Haidar, R. Akl, and Y. Chan. Enhanced Channel Assignment and Load Distribution in IEEE 802.11 WLANs. In *IEEE International Conference on Signal Processing and Communications, ICSPC'07*, pages 768–771, November 2007.

[49] B.E. Alexander, editor. *802.11 Wireless Network Site Surveying and Installation.* Cisco Press, November 2004.

[50] C. Andren and M. Webster. CCK Modulation Delivers 11 Mbps for High Rate IEEE 802.11 Extension. In *Proc. Wireless Symposium/Portable by Design Conference*, 1999.

[51] C. Andren, M. Webster, and K. Halford. CCK: the new IEEE 802.11 standard for 2.4 GHZ wireless LANs. In *Proceedings of International IC-Taipei Conference*, pages 25–39, May 2000.

[52] V. Angelakis, S. Papadakis, V. Siris, and A. Traganitis. Adjacent channel interference in 802.11a: Modeling and testbed validation. In *IEEE Radio and Wireless Symposium*, pages 591–594, January 2008.

[53] T. Arildsen and FHP. Fitzek. The C-Cube Concept - Combining Cross-Layer Protocol Design, Cognitive-, and Cooperative Network Concepts. In F. Fittzek and M. Katz, editors, *Cognitive Wireless Networks*, chapter 21. Springer, 2007.

[54] P. Bahl, M.T. Hajiaghayi, K. Jain, S.V. Mirrokni, L. Qiu, and A. Saberi. Cell Breathing in Wireless LANs: Algorithms and Evaluation. *IEEE Transactions on Mobile Computing*, 6(2):164–178, 2007.

[55] Anand Balachandran, Paramvir Bahl, and Geoffrey M. Voelker. Hot-Spot Congestion Relief in Public-area Wireless Networks. In *Proc. of 4th IEEE Workshop on Mobile Computing Systems and Applications*, pages 70–80, June 2002.

[56] Anand Balachandran, Geoffrey M. Voelker, Paramvir Bahl, and P. Venkat Rangan. Characterizing user behavior and network performance in a public wireless lan. In *Proceedings of ACM SIGMETRICS 2002*, volume 30, pages 195–205, June 2002.

[57] Magdalena Balazinska and Paul Castro. Characterizing mobility and network usage in a corporate wireless local-area network. In *Proc. of First International Conference on Mobile Systems, Applications, and Services, MobiSys'03*, May 2003.

[58] N. Baldo, F. Maguolo, and M. Miozzo. A new approach to simulating PHY, MAC and Routing. In *Proceeding from the 2008 workshop on ns-2: the IP network simulator, WNS2'08*, page 12, October 2008.

[59] E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa. On the Performance of IEEE 802.11 under Jamming. In *The 27th Conference on Computer Communications, INFOCOM'08*, pages 1265–1273, April 2008.

[60] Alessandro Bazzi, Marco Diolaiti, and Gianni Pasolini. Measurement based Call Admission Control Strategies in Infrastructured IEEE 802.11. In *The 16th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC 2005.*, September 2005.

[61] Alessandro Bazzi, Marco Diolaiti, and Gianni Pasolini. Link adaptation algorithms over IEEE 802.11 WLANs in collision prone channels. In *IEEE 63rd Vehicular Technology Conference, VTC'06-Spring*, May 2006.

[62] Yigal Bejerano and Seung-Jae Han. Cell Breathing Techniques for Load Balancing in Wireless LANs. In *25th Annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM'06*, pages 1–13, April 2006.

[63] Yigal Bejerano, Seung-Jae Han, and Li (Erran) Li. Fairness and load balancing in wireless LANs using association control. In *Proc. of the 10th international conference on Mobile computing and networking, MobiCom'04*, pages 315— 329, 2004.

[64] Albert Beltran Roman. Punts d'accés 802.11 intel·ligents basats en linux (I). Master's thesis, Escola Politècnica Superior de Castelldefels, EPSC - UPC, February 2005.

[65] M.G. Di Benedetto and L. De Nardis. Tuning UWB signals by pulse shaping. *EURASIP Journal on Signal Processing, Special Issue on Signal Processing in UWB Communications*, 2005.

[66] M. Bernaschi, F. Cacace, G. Iannello, M. Vellucci, and L. Vollero. OpenCAPWAP: An open source CAPWAP implementation for the management and configuration of WiFi hot-spots. *Computer Networks*, 53(2):217–230, 2009.

[67] RA. Berry and EM. Yeh. Cross-Layer Wireless Resource Allocation. *IEEE Signal Processing magazine*, pages 59–68, 2004.

[68] Vaduvur Bharghavan, Alan Demers, Scott Shenker, and Lixia Zhang. MACAW: a media access protocol for wireless LAN's. In *Proceedings of the conference on Communications architectures, protocols and applications, SIGCOMM'94*, pages 212–225, New York, NY, USA, 1994. ACM.

[69] C. Bianchi and A. Meloni. Natural and man-made terrestrial electromagnetic noise: an outlook. *Annals og Geophysics*, 50(3):435–445, June 2007.

[70] G. Bianchi. Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, March 2000.

[71] G. Bianchi and I. Tinnirello. Kalman filter estimation of the number of competing terminals in an IEEE 802.11 network. In *22nd Annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM'03*, volume 2, pages 844–852, April 2003.

[72] Giuseppe Bianchi and Ilenia Tinnirello. Improving Load Balancing Mechanisms in Wireless Packet Networks. In *IEEE International Conference on Communications 2002, ICC'02*, volume 2, pages 891–895, April 2002.

[73] Saad Biaz and Shaoen Wu. Rate Adaptation Algorithms for IEEE 802.11 Networks: A Survey and Comparison. In *13th IEEE Symposium on Computers and Communications, ISCC'08*, July 2008.

[74] Nikita Borisov, Ian Goldberg, and David Wagner. Intercepting mobile communications: the insecurity of 802.11. In *MobiCom'01: Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 180–189. ACM, 2001.

[75] P. O. Börjesson and C. E. Sundberg. Simple approximations of the error function Q(x) for communications applications. *IEEE Transactions on Communications*, 27(1):639–643, March 1979.

[76] R. Borndörfer, A. Eisenblätter, M. Grötschel, and A. Martin. Frequency Assignment in Cellular Phone Networks. *Annals of Operations Research*, 76:73–93, 1998.

[77] D. Brélaz. New Methods to Color the Vertices of a Graph. *Communications of the ACM*, 22:251–256, 1979.

[78] Olivia Brickley, Susan Rea, and Dirk Pesch. Load Balancing for QoS Optimisation in Wireless LANs Utilising Advanced Cell Breathing Techniques. In *IEEE 61st Vehicular Technology Conference. VTC'05-Spring*, May 2005.

[79] R. Bruno, M. Conti, and E. Gregori. IEEE 802.11 optimal performances: RTS/CTS mechanism vs. basic access. In *13th IEEE International Symposium On Personal, Indoor And Mobile Radio Communications, PIMRC'02*, volume 4, pages 1747–1751, September 2002.

[80] P. Calhoun, M. Montemurro, and D. Stanley. Control and Provisioning of Wireless Access Points (CAPWAP) Protocol Binding for IEEE 802.11. RFC 5416 (Standards Track), March 2009.

[81] P. Calhoun, M. Montemurro, and D. Stanley. Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification. RFC 5415 (Standards Track), March 2009.

[82] P. Calhoun, B. O'Hara, R. Suri, N. Cam Winget, S. Kelly, M. Williams, and S. Hares. Light Weight Access Point Protocol (LWAPP). draft-ohara-capwap-lwapp-04, March 2007.

[83] F. Cali, M. Conti, and E. Gregori. Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit. *IEEE/ACM Transactions on Networking (TON)*, 8(6):785–799, December 2000.

[84] Gion Reto Cantieni, Qiang Ni, Chadi Barakat, and Thierry Turletti. Performance analysis under finite load and improvements for multirate 802.11. *Computer Communications Journal, special issue on Performance Issues of Wireless LANs, PANs, and Ad Hoc Networks*, 28(10):1095–1109, June 2005.

[85] M.M. Carvalho and J.J. García-Luna-Aceves. Delay analysis of IEEE 802.11 in single-hop networks. In *11th IEEE International Conference on Network Protocols, ICNP'03*, pages 146–155, November 2003.

[86] Antonio Carzaniga, Gian Pietro Picco, and Giovanni Vigna. Designing distributed applications with a mobile code paradigm. In *Proceedings of the 19th International Conference on Software Engineering*, Boston, MA, USA, 1997.

[87] J. Case, R. Mundy, D. Partain, and B. Stewart. Introduction to Version 3 of the Internet-standard Network Management Framework. RFC 2570 (Informational), April 1999.

[88] J.D. Case, M. Fedor, M.L. Schoffstall, and J. Davin. Simple Network Management Protocol (SNMP). RFC 1157 (Historic), May 1990.

[89] M. Chandra and A. Roy. Extensions to OSPF to Support Mobile Ad Hoc Networking. ddraft-ietf-ospf-manet-or-01, September 2008.

[90] P. Chatzimisios, A.C. Boucouvalas, and V. Vistas. Influence of channel BER on IEEE 802.11 DCF. *Electronics Letters*, 39(23):1687–1689, November 2003.

[91] P. Chatzimisios, A.C. Boucouvalas, and V. Vistas. Packet delay analysis of IEEE 802.11 MAC protocol. *Electronics Letters*, 39(18):1358–1359, September 2003.

[92] C. Chaudet, D. Dhoutaut, and I.G. Lassous. Experiments of some performance issues with IEEE 802.11b in ad hoc networks. In *2nd Annual Conference on Wireless On-Demand Network Systems and Services, WONS'05*, January 2005.

[93] C. Chaudet, D. Dhoutaut, and I.G. Lassous. Performance issues with IEEE 802.11 in Ad Hoc Networking. *IEEE Communications Magazine*, 43(7):110–116, July 2005.

[94] Morsy M. Cheikhrouhou, Pierre Conti, and Jacques Labetoulle. Intelligent agents in network management, a state-of-the-art. *Networking and Information Systems*, 1(1):9–38, 1998.

[95] T. Chen, H. Zhang, GM. Maggio, and I. Chlamtac. CogMesh: A Cluster-Based Cognitive Radio Network. In *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, DySPAN'07*, pages 168–178, April 2007.

[96] Thomas M. Chen and Stephen S. Liu. A Model and Evaluation of Distributed Network Management Approaches. *IEEE Journal on Selected Areas in Communications*, 20(4):850–857, May 2002.

[97] Walter Y. Chen. *Home Networking Basis: Transmission Environments and Wired/Wireless Protocols*. Prentice Hall, July 2003.

[98] H.S. Chhaya and S. Gupta. Performance Modeling of Asynchronous Data Transfer Methods in the IEEE 802.11 MAC Protocol. *Wireless Networks*, 3:217–234, 1997.

[99] C.F. Chiasserini and R.R. Rao. Coexistence mechanisms for interference mitigation in the 2.4-ghz ism band. *IEEE Transactions on Wireless Communications*, 2(5):964–975, September 2003.

[100] I. Chlamtac, M. Conti, and J.J-N Liu. Mobile ad hoc networking: imperatives and challenges. *Ad Hoc Networks*, 1(1):13–64, July 2003.

[101] Sunwoong Choi, Kihong Park, and Chong kwon Kim. On the performance characteristics of WLANs: revisited. In *Proceedings of the ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 97–108. ACM, 2005.

[102] T. Clausen and P. Jacquet. Optimized Link State Routing Protocol (OLSR). RFC 3626 (Experimental), October 2003.

[103] D. Costa. On the Use of some Known Methods for T-Colourings of Graphs. *Annals of Operations Research*, 41:343–358, 1993.

[104] J. del Prado and S. Choi. Link Adaptation Strategy for IEEE 802.11 WLAN via Received Signal Strength Measurement. In *IEEE International Conference on Communications, 2003. ICC '03*, volume 2, pages 1108–1113, May 2003.

[105] Imed Ben Dhou. A Novel Load-Sharing Algorithm for Energy Efficient MAC Protocol Compliant with 802.11 WLAN. In *IEEE 50th Vehicular Technology Conference. VTC'99-Fall*, volume 2, pages 1238–1242, September 1999.

[106] R. Draves, J. Padhye, and B. Zill. Routing in Multi-radio, Multi-hop Wireless Mesh Networks. In *Proc. of the 10th international conference on Mobile computing and networking, MobiCom'04*, September 2004.

[107] M. Drieberg, F-C. Zheng, R. Ahmad, and S. Olafsson. An Asynchronous Distributed Dynamic Channel Assignment Scheme for Dense WLANs. In *IEEE International Conference on Communications, ICC'08*, pages 2507–2511, May 2008.

[108] A. Eisenblätter, H-M. Geerdes, and I. Siomina. Integrated Access Point Placement and Channel Assignment for Wireless LANs in an Indoor Office Environment. In *8th IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks, WoWMoM'07.*, June 2007.

[109] A. Eisenblätter, M. Grötschel, and A. M. C. A. Koster. Frequency planning and ramifications of coloring. *Discussiones Mathematicae Graph Theory*, 22(1):51–88, 2002.

[110] O. Ekici and A. Yongacoglu. A novel association algorithm for congestion relief in IEEE 802.11 WLANs. In *Proceedings of the international conference on Wireless communications and mobile computing, IWCMC'06*, pages 725–730, 2006.

[111] O. Ekici and A. Yongacoglu. Balanced association algorithm for IEEE 802.11 extended service areas. *Wireless Communications and Mobile Computing*, ??(??):??, ?? 2008.

[112] Wassim El-Hajj and Hamed Alazemi. Optimal frequency assignment for IEEE 802.11 wireless networks. *Wireless Communications and Mobile Computing*, 9(1):131–141, January 2009.

[113] M. Ergen, B. Dundar, and P. Varaiya. Throughput analysis of an extended service set in IEEE 802.11. In *Global Telecommunications Conference, GLOBECOM'04*, volume 2, pages 1040–1045, November 2004.

[114] M. Ergen and P. Varaiya. Throughput Analysis and Admission Control for IEEE 802.11a. *Mobile Networks and Applications*, 10(5):705–716, October 2005.

[115] M. Lacage et. al. Yet Another Network Simulator (YANS), v. 0.90.
http://yans.inria.fr/, May 2006.

[116] Lluís Faixó. Redes mesh basadas en puntos de acceso inteligentes 802.11 open source (I). Master's thesis, Escola Politècnica Superior de Castelldefels, EPSC - UPC, September 2005.

[117] Marco Fiore. ns-2 wireless update patch.
http://www.tlc-networks.polito.it/fiore/, 2004.

[118] M. Fleury, G. Flores, and M. Reed. Clarification of the "OPNET NS-2 Comparison" Paper with regards to OPNET Modeler.
http://privatewww.essex.ac.uk/˜fleum/OPNET-NS2_Comparison.pdf, 2004.

[119] G. Flores, M. Paredes, E. Jammeh, M. Fleury, and M.J. Reed. OPNET modeler and Ns-2 - Comparing the accuracy of network simulators for packet-level analysis using a network testbed. *WSEAS Transactions on Computers*, 2(3):700–707, July 2003.

[120] Scott R. Fluhrer, Itsik Mantin, and Adi Shamir. Weaknesses in the Key Scheduling Algorithm of RC4. In *SAC01: Revised Papers from the 8th Annual International Workshop on Selected Areas in Cryptography*, pages 1–24. Springer-Verlag, 2001.

[121] S. Frattasi, H. Fathi, FHP. Fitzek, R. Prasad, and MD. Katz. Defining 4G technology from the user's perspective. *IEEE Network*, 20(1):35–41, February 2006.

[122] Y. Fukuda, T. Abe, and Y. Oie. Decentralized Access Point Selection Architecture for Wireless LANs. In *Proceedings of the Wireless Telecommunications Symposium, WTS'04*, pages 137–145, September 2004.

[123] Y. Fukuda, M. Honjo, and Y. Oie. Development of Access Point Selection Architecture with Avoiding Interference for WLANs. In *The 17th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC'06.*, September 2006.

[124] Y. Fukuda and Y. Oie. Decentralized Access Point Selection Architecture for Wireless LANs. *IECE Transactions on Communications*, E90-B(9):2513–2523, September 2007.

[125] Fiorenzo Gamba, Jean-Frédéric Wagen, and Daniel Rossier. Towards adaptive wlan frequency management using intelligent agents. In *Ad-Hoc, Mobile, and Wireless Networks, Second International Conference, ADHOC-NOW 2003 Montreal*, pages 116–127, 2003.

[126] E. Garcia. Available Admission Capacity Estimations in IEEE 802.11 Access Points. Tech. report available online at
http://hdl.handle.net/2117/2045, UPC, April 2008.

[127] E. Garcia, Lluís Faixó, Rafael Vidal, and Josep Paradells. Inter-access point communications for distributed resource management in 802.11 networks. In *4th international workshop on Wireless mobile applications and services on WLAN hotspots, WMASH'06*, pages 11–19, New York, NY, USA, September 2006. ACM Press.

[128] E. Garcia, J.L. Ferrer, E. López, R. Vidal, and J. Paradells. Client-driven load balancing through association control in IEEE 802.11 WLANs. *European Transactions on Telecommunications*, 20(5):494–507, August 2009.

[129] E. Garcia, E. López, R. Vidal, and J. Paradells. Effect of adjacent-channel interference in IEEE 802.11 WLANs. In *2nd Int. Conference on Cognitive Radio Oriented Wireless Networks and Communications, CrownCom'07*, pages 118–125, August 2007.

[130] E. Garcia, E. López, R. Vidal, and J. Paradells. IEEE Wireless LAN capacity in multicell environments with rate adaptation. In *The 18th IEEE International*

Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC'07., September 2007.

[131] E. Garcia, D. Viamonte, R. Vidal, and J. Paradells. Achievable Bandwidth Estimation for Stations in Multi-Rate IEEE 802.11 WLAN Cells. In *8th IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks, WoWMoM'07.*, pages 1–8, June 2007.

[132] E. Garcia and R. Vidal. Localización en redes WLAN 802.11: desarrollo e implementación de una solución basada en traps SNMP. In *IV Jornadas de Ingeniería Telemática, JITEL 2003*, September 2003.

[133] E. Garcia, R. Vidal, and J. Paradells. Implementation of a distributed dynamic channel assignment mechanism for IEEE 802.11 Networks. In *The 16th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC 2005.*, September 2005.

[134] E. Garcia, R. Vidal, and J. Paradells. New Algorithm for Distributed Frequency Assignments in IEEE 802.11 Wireless Networks. In *11th European Wireless Conference 2005, EW05*, volume 1, pages 211–217, April 2005.

[135] E. Garcia, R. Vidal, and J. Paradells. Capacity Prediction in Ad-hoc IEEE 802.11 WLANs. In *(poster) ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc'06.*, May 2006.

[136] E. Garcia, R. Vidal, and J. Paradells. Load Balancing in WLANs through IEEE 802.11k Mechanisms. In *11th IEEE Symposium on Computers and Communications, ISCC'06.*, June 2006.

[137] E. Garcia, R. Vidal, and J. Paradells. Cooperative Load Balancing in IEEE 802.11 Networks with Cell Breathing. In *13th IEEE Symposium on Computers and Communications, ISCC'08.*, pages 1133–1140, July 2008.

[138] E. Garcia, R. Vidal, and J. Paradells. Frequency assignments in IEEE 802.11 WLANs with efficient spectrum sharing. *Wireless Communications and Mobile Computing*, 9(8):1125–1140, August 2009.

[139] Víctor N. García Fernández. Redes mesh basadas en puntos de acceso inteligentes 802.11 open source (II). Master's thesis, Escola Politècnica Superior de Castelldefels, EPSC - UPC, September 2005.

[140] M. Garetto, T. Salonidis, and E.W. Knightly. Modeling per-flow throughput and capturing starvation in csma multi-hop wireless networks. In *Proceedings of the 25th IEEE Annual Conference INFOCOM'06*, April 2006.

[141] M. Garetto, T. Salonidis, and E.W. Knightly. Modeling Per-Flow Throughput and Capturing Starvation in CSMA Multi-Hop Wireless Networks. *IEEE/ACM Transactions on Networking*, 16(4):864–877, August 2008.

[142] Matthew Gast. *802.11 Wireless Networks: The Definitive Guide*. O'Reilly, 2nd edition, April 2005.

[143] P. Gawthrop, F. Sanders, K. Nebbia, and J. Sell. Radio Spectrum Measurements of Individual Microwave Ovens. Report on Emission Spectrum Measurements on Individual Transmitters 94-303-1, NTIA, March 1994.

[144] P. Gawthrop, F. Sanders, K. Nebbia, and J. Sell. Radio Spectrum Measurements of Individual Microwave Ovens. Report on Emission Spectrum Measurements on Individual Transmitters 94-303-2, NTIA, March 1994.

[145] J. Geier. *Wireless Networks First-Step*. Cisco Press, August 2004.

[146] Moisés Gómez Díaz. Detecció d'atacs DoS amb inhibidors de freqüències sobre xarxes IEEE 802.11. Master's thesis, Escola Politècnica Superior de Castelldefels, EPSC - UPC, May 2009.

[147] Huazhi Gong and JongWon Kim. Dynamic load balancing through association control of mobile users in WiFi networks. *IEEE Transactions on Consumer Electronics*, 54(2):342–348, May 2008.

[148] Fabrizo Granelli, Honggang Zhang, Xiaofei Zhou, and Stefano Maranò. Research advances in cognitive ultra wide band radio and their application to sensor networks. *Mobile Networks Applications*, 11(4):487–499, 2006.

[149] D.B. Green and M.S. Obaidat. An Accurate Line of Sight Performance Model for Ad-Hoc 802.11 Wireless LAN (WLAN) Devices. In *IEEE International Conference on Communications 2002, ICC 2002*, volume 5, pages 3424–3428, 2002.

[150] Ramakrishna Gummadi, David Wetherall, Ben Greenstein, and Srinivasan Seshan. Understanding and mitigating the impact of RF interference on 802.11 networks. In *Applications, technologies, architectures, and protocols for computer communications, SIGCOMM'07*, pages 385–396. ACM, 2007.

[151] D. Haccoun and G. Begin. High-rate punctured convolutional codes for viterbi and sequential decoding. *IEEE Transactions on Communications*, 37(11):1113–1125, November 1989.

[152] K. Halford, S. Halford, M. Webster, and C. Andren. Complementary Code Keying for RAKE-Based Indoor Wireless Communications. In *IEEE IEEE International Symposium on Circuits and Systems, ISCAS'99*, volume 4, pages 427–430, 1999.

[153] D. Harrington, R. Presuhn, and B. Wijnen. An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks. RFC 3411 (Standards Track), Decemberaug 2002.

[154] M. Hata. Empirical formula for propagation loss in land mobile radio services. *IEEE Transactions on Vehicular Technology*, 29(3):317–325, August 1980.

[155] Simon Haykin. *Communications Systems*. John Wiley and Sons, New York, NY, USA, 4th edition, May 2000.

[156] Simon Haykin. Cognitive radio: brain-empowered wireless communications. *IEEE Journal on Selected Areas in Communications*, 23:317–325, February 2005.

[157] Klaus Heck. Wireless LAN performance in overlapping cells. In *58th Vehicular Technology Conference, IEEE VTC'03-Fall*, volume 5, pages 2895–2900, October 2003.

[158] Tristan Henderson, David Kotz, and Ilya Abyzov. The changing usage of a mature campus-wide wireless network. In *Proc. of the 10th annual international conference on Mobile computing and networking*, pages 187–201, September 2004.

[159] Benjamin E. Henty. A Brief Tutorial on the PHY and MAC layers of the IEEE 802.11b Standard. White paper, Intersil, July 2001.

[160] Martin Heusse, Franck Rousseau, Gilles Berger-Sabbatel, and Andrzej Duda. Performance anomaly of 802.11b. In *Proceedings of the 22nd IEEE Annual Conference INFOCOM'03*, volume 2, pages 836–843, March 2003.

[161] A. Hills. Large-Scale Wireless LAN Design. *IEEE Communications Magazine*, 39(11):98–107, November 2001.

[162] A. Hills and B. Friday. Radio resource management in wireless LANs. *IEEE Radio Communications Magazine*, 42(12):S9–14, December 2004.

[163] G. Holland, N. Vaidya, and P. Bahl. A rate-adaptive mac protocol for multi-hop wireless networks. In *Proceedings of 7th International Conference on Mobile Computing and Networks, ACM MOBICOM 2001*, pages 236–251, 2001.

[164] TC Hou, LF Tsao, and HC Liu. Analyzing the Throughput of IEEE 802.11 DCF Scheme with Hidden Nodes. In *58th IEEE Vehicular Technology Conference, VTC-Fall'03*, volume 5, pages 2870–2874, October 2003.

[165] TC Hou, LF Tsao, and HC Liu. Throghput Analysis of the IEEE 802.11 DCF Scheme in Multi-hop Ad-hoc Networks. In *International Conference on Wireless Networks, ICWN'03*, June 2003.

[166] J. Li and C. Blake and D. De Couto and H.I. Lee and M. Morris. Capacity of Ad Hoc wireless networks. In *Proceedings of the 7th annual international conference on Mobile computing and networking, MobiCom'01*, pages 61–69. ACM, 2001.

[167] Issam Jabri, Nicolas Krommenacker, Thierry Divoux, and Adel Soudani. IEEE 802.11 Load Balancing: An Approach for QoS Enhancement. *International Journal of Wireless Information Networks*, 15(1):16–30, March 2008.

[168] R. Jain, D. Chiu, and W. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Technical Report TR-301, DEC Research, September 1984.

[169] B-S. Jeon, E-J. Ko, and G-H. Lee. Network Management System for Wireless LAN Service. In *10th International Conference on Telecommunications, 2003. ICT 2003*, volume 2, pages 948–953, March 2003.

[170] J.R.R. Tolkien. *The Fellowship of the Ring*. Allen & Unwin, 1954.

[171] Huei-Jiun Ju, I. Rubin, and Yen-Chang Kuan. An adaptive RTS/CTS control mechanism for IEEE 802.11 MAC protocol. In *IEEE 57th Vehicular Technology Conference, VTC'03-Spring*, volume 2, pages 1469–1473, April 2003.

[172] J. Jun, P. Peddabachagari, and M Sichitiu. Theoretical maximum Throughput of IEEE 802.11 and its Applications. In *Proceedings of the Second IEEE International Symposium on Network Computing and Applications*, page 249, April 2003.

[173] M. Kahani and H. Beadle. Decentralised approaches for network management. *ACM SIGCOMM Computer Communication Review*, 27(3):36–47, July 1997.

[174] A. Kamerman and N. Erkocevic. Microwave oven interference on wireless LANs operating in the 2.4 GHz ISM band. In *The 8th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications PIMRC'97*, volume 3, pages 1221–1227, September 1997.

[175] A. Kamerman and L. Monteban. WaveLAN II: A high-performance wireless LAN for the unlicensed band . *Bell Labs Techical Journal*, pages 118–133, 1997.

[176] Kameswari Chebrolu and Bhaskaran Raman and Sayandeep Sen. Long-distance 802.11b links: performance measurements and experience. In *Proceedings of the 12th annual international conference on Mobile computing and networking, MobiCom'06*, pages 74–85. ACM, 2006.

[177] J. Kantorocitch and P. Mähönen. Case studies and experiments of SNMP in wireless networks. In *IEEE Workshop on IP Operations and Management, 2002*, pages 179–183, 2002.

[178] Anand Kashyap, Samrat Ganguly, and Samir R. Das. A measurement-based approach to modeling link capacity in 802.11-based wireless networks. In *Proceedings of the 13th annual ACM international conference on Mobile computing and networking, MobiCom'07*, pages 242–253, New York, NY, USA, 2007. ACM.

[179] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot. Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks. In *26th IEEE International Conference on Computer Communications, INFOCOM'07*, pages 1451–1459, May 2007.

[180] S. Keshav. The real network simulator, v5.0.
http://www.cs.cornell.edu/skeshav/real/, 1997.

[181] Masood Khosroshahy. Study and Implementation of IEEE 802.11 Physical Channel Model in YANS (NS3 prototype) Network Simulator. Technical report, INRIA-Sophia Antipolis-Planète Group, November 2006.

[182] Masood Khosroshahy, Thierry Turletti, and Katia Obraczka. Snapshot of MAC, PHY and Propagation Models for IEEE 802.11 in Open-Source Network Simulators. Technical report, INRIA-Sophia Antipolis-Planète Group, September 2007.

[183] Ki Hong Kim, Jung Ha Kim, Young Joong Yoon, Jae Ho Seok, and Jae Woo Lim. Propagation model for the WLAN service at the campus environments. In *The 57th IEEE Semiannual Vehicular Technology Conference, VTC'03-Spring.*, volume 1, pages 196–200, April 2003.

[184] Y-J. Kim and Y-J. Suh. An Efficient Rate Switching Scheme for IEEE 802.11 Wireless LANs. In *IEEE 61st Vehicular Technology Conference, VTC'05-Spring*, May 2005.

[185] K.K. Leung and B. McNair and L.J. Cimini and J.H. Winters. Outdoor IEEE 802.11 cellular networks: MAC protocol design and performance. In *IEEE International Conference on Communications, ICC'02.*, volume 1, pages 595–599, 2002.

[186] W. Klotz. Graph coloring algorithms. Mathematical report, TU-Clausthal, May 2002.

[187] D. Kotz, C. Newport, and C. Elliott. The mistaken axioms of wireless-network research. Technical report, Dept. of Computer Science, Dartmouth College, July 2003.

[188] D. Kotz, C. Newport, R.S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *Proceedings of the 7th ACM MSWiM*, pages 78–82, 2004.

[189] A. Kumar, E. Altman, D. Miorandi, and M. Goyal. New Insights from a Fixed Point Analysis of Single Cell IEEE 802.11 WLANs. In *24th IEEE International Conference on Computer Communications, INFOCOM'05*, volume 3, pages 1550–1561, March 2005.

[190] A. Kumar and V. Kumar. Optimal Association of Stations and APs in an IEEE 802.11 WLAN. In *National Conference on Communications, 2005. NCC 2005.*, January 2005.

[191] A. Kumar and S. Roy. The Enhanced Network Simulator TeNS. http://www.cse.iitk.ac.in/users/braman/tens/.

[192] V. Kumar and R.V.R. Kumar. Performance analysis of a finite word length implemented CCK modem with rake receiver for WLAN system. In *1st IEEE and IFIP International Conference on Internet in Central Asia*, September 2005.

[193] Mathieu Lacage and Thomas R. Henderson. Yet another network simulator. In *WNS2 '06: Proceeding from the 2006 workshop on ns-2: the IP network simulator*, page 12, New York, NY, USA, 2006. ACM.

[194] Mathieu Lacage, Mohammad Hossein Manshaei, and Thierry Turletti. IEEE 802.11 Rate Adaptation: A Practical Approach. In *Proceedings of the 7th ACM MSWiM*, pages 126–134, October 2004.

[195] Heeyoung Lee, Seongkwan Kim, Okhwan Lee, Sunghyun Choi, and Sung-Ju Lee. Available bandwidth-based association in IEEE 802.11 Wireless LANs. In *Proceedings of the 11th international symposium on Modeling, analysis and simulation of wireless and mobile systems, MSWiM'08*, pages 132–139, New York, NY, USA, 2008. ACM.

[196] William C.Y. Lee. *Mobile Communications Design Fundamentals.* John Wiley & Sons, second edition, February 1993.

[197] Y. Lee, K. Kim, and Y. Choi. Optimization of AP Placement and Channel Assignment in Wireless LANs. In *Proceedings of the 27th Annual IEEE Conference on Local Computer Networks*, pages 831–836, November 2002.

[198] Sam Leffler. The MadWifi project. http://madwifi-project.org, March 2009.

[199] K.K. Leung and B-J. Kim. Frequency assignment for IEEE 802.11 wireless networks. In *IEEE 58th Vehicular Technology Conference, VTC'03-Fall*, volume 3, pages 1422–1426, October 2003.

[200] A. Levanti, F. Giordano, and I. Tinnirello. A CAPWAP Architecture for Automatic Frequency Planning in WLAN. In *12th IEEE Symposium on Computers and Communications, ISCC'07*, July 2007.

[201] D. Levi and J. Schoenwaelder. Definitions of Managed Objects for the Delegation of Management Scripts. RFC 3165 (Proposed Standard), August 2001.

[202] Yan Li, Xiaowen Wang, and S.A. Mujtaba. Co-channel interference avoidance algorithm in 802.11 wireless LANs. In *IEEE 58th Vehicular Technology Conference, VTC'03-Fall*, volume 4, pages 2610–2614, October 2003.

[203] Xiang Ling and K.L. Yeung. Joint access point placement and channel assignment for 802.11 wireless LANs. In *IEEE Wireless Communications and Networking Conference, 2005. WCNC05*, volume 3, pages 1583–1588, March 2005.

[204] C. Liu and F. Lee. Spectrum modelling of OFDM signals for WLAN. *Electronics Letters*, 40(22):1431–1432, October 2004.

[205] Shao-bo Liu, Aiping Huang, Zhao-yang Zhang, and Zhijian Zhang. Performance analysis of cck modulation under multipath fading channel. In *Proceedings of the 6th Nordic Signal Processing Symposium, NORSIG'04*, pages 276–279, June 2004.

[206] A. Longley and P. Rice. Prediction of tropospheric radio transmission over irregular terrain, a computer method. Tech. Rep. ERL 79-ITS 67, ESSA, Washington DC, July 1968.

[207] Luís Daniel Ruiz López and Marc Oliveras Pla. Redes mesh basadas en puntos de acceso inteligentes 802.11 open source (III). Master's thesis, Escola Politècnica Superior de Castelldefels, EPSC - UPC, February 2006.

[208] E. López-Aguilera. *Contributions to the Evaluation and Enhancement of WLAN IEEE 802.11 Medium Access Control Mechanism.* PhD thesis, Technical University of Catalonia (UPC), June 2008.

[209] E. López-Aguilera, J. Casademont, and J. Cotrina. Outdoor IEEE 802.11g Cellular Network Performance. In *Global Telecommunications Conference, GLOBECOM'04*, volume 5, pages 2992–2996, November 2004.

[210] E. López-Aguilera, J. Casademont, and E. Garcia. A study on the Influence of Transmission Errors on WLAN IEEE 802.11 MAC Performance. *Wireless Communications and Mobile Computing*, 2009.

[211] Hui Luo and N.K. Shankaranarayanan. A distributed dynamic channel allocation technique for throughput improvement in a dense wlan environment. In *Acoustics, Speech, and Signal Processing, 2004. Proceedings. (ICASSP '04)*, volume 5, pages 345–348, May 2004.

[212] Yao Ma, Teng Loon Lim, and S. Pasupathy. Error probability for coherent and differential PSK over arbitrary Rician fading channels with multiple cochannel interferers. *IEEE Transactions on Communications*, 50(3):429–441, March 2002.

[213] P. Mähönen and et. al. The ARAGORN Project: Adaptive Reconfigurable Access and Generic interfaces for Optimisation in Radio Networks. Technical report, 2007.

[214] P. Mähönen and et. al. ARAGORN - State of the art. Technical Report D 2.1, April 2008.

[215] P. Mähönen, J. Riihijärvi, and M. Petrova. Automatic channel allocation for small wireles local area networks using graph colouring algorithm approach. In *15th IEEE International Symposium On Personal, Indoor And Mobile Radio Communications, PIMRC'04, Barcelona*, September 2004.

[216] WQ. Malik and DJ. Edwards. Cognitive Techniques for Ultrawideband Communications. In *Seminar on Ultra Wideband Systems, Technologies and Applications*, pages 81–86, April 2006.

[217] Jouni Malinen. Host AP driver for Intersil Prism2/2.5/3, hostapd, and WPA Supplicant. http://hostap.epitest.fi, March 2009.

[218] Stefan Mangold and Lars Berlemann. IEEE 802.11k: Improving Confidence in Radio Resource Measurements. In *The 16th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC'05.*, September 2005.

[219] S. Mccanne, S. Floyd, and K. Fall et. al. ns-2 (Network Simulator 2), Version 2.33. http://www.isi.edu/nsnam/ns/, March 2008.

[220] K. McCloghrie and M.T. Rose. Management Information Base for Network Management of TCP/IP-based internets:MIB-II. RFC 1213 (Standard), March 1991. Updated by RFCs 2011, 2012, 2013.

[221] K. Medepalli and F.A. Tobagi. Throughput analysis of IEEE 802.11 wireless LANs using an average cycle time approach. In *IEEE Global Telecommunications Conference, GLOBECOM'05*, pages 3007–3011, November 2005.

[222] Arunesh Mishra, Suman Banerjee, and William Arbaugh. Weighted Coloring based Channel Assignment for WLANs. *Mobile Computing and Communications Review*, 9(3):19–31, July 2005.

[223] Arunesh Mishra, V. Brick, A. Srinivasan, and W. Arbaugh W. A client-driven approach for channel management in wireless lans. In *25th Conference on Computer Communications, IEEE Infocom'06*, April 2006.

[224] Arunesh Mishra, Vivek Shrivastava, Dheeraj Agrawal, Suman Banerjee, and Samrat Ganguly. Distributed channel management in uncoordinated wireless environments. In *Proceedings of the 12th annual international conference on Mobile computing and networking, MobiCom'06*, pages 170–181, New York, NY, USA, 2006. ACM.

[225] Arunesh Mishra, Vivek Shrivastava, Suman Banerjee, and William Arbaugh. Partially overlapped channels not considered harmful. *SIGMETRICS Perform. Eval. Rev.*, 34(1):63–74, 2006.

[226] J. Mitola. The software radio architecture. *IEEE Communications Magazine*, 33:26–33, May 1995.

[227] J. Mitola and C. Maguire. Cognitive radio: Making software radios more personal. *IEEE Personal Communications Magazine*, 6:13–18, August 1999.

[228] S. Miyamoto and N. Morinaga. Effect of microwave oven interference on the performance of digital radio communications systems. In *IEEE International Conference on Communications 1997, ICC'97*, pages 51–55, June 1997.

[229] S. Miyamoto, Y. Yamanaka, and T. Shinozuka. A study of the effect of microwave oven interference on the performance of digital radio communications systems. *Electron. Commun. Japan*, 80(12):58–67, 1997.

[230] P. Mogensen, P. Jensen, and J. Andersen. 1800 mhz mobile net planning based on 900mhz measurements. COST 231 TD (91)-08, European Comission - COST, Firenze, January 1991.

[231] S. Muthuswamy, I. Marsic, and A. Annamalai. New methods for estimating/forecasting link bandwidths in 802.11b WLANs. In *IEEE 60th Vehicular Technology Conference, VTC'04-Fall*, volume 2, pages 1163–1168, September 2004.

[232] M.V. Clark and K.K Leung and B. McNair and Z. Kostic. Outdoor IEEE 802.11 cellular networks: radio link performance. In *IEEE International Conference on Communications, ICC'02.*, volume 1, pages 512–516, 2002.

[233] V. Navda, A. Bohra, S. Ganguly, and D. Rubenstein. Using channel hopping to increase 802.11 resilience to jamming attacks. In *26th IEEE International Conference on Computer Communications, INFOCOM'07*, pages 2526–2530, May 2007.

[234] P.R. Nelson. An approximation for the complex normal probability integral. *BIT Numerical Mathematics*, 22(1):79–129, March 1982.

[235] Quoc-Thinh Nguyen-Vuong, N. Agoulmine, and Y. Ghamri-Doudane. A user-centric and context-aware solution to interface management and access network selection in heterogeneous wireless environments. *Computer Networks*, 52(18):3358–3372, December 2008.

[236] Q. Ni, I. Aad, C. Barakat, and T. Turletti. Modeling and Analysis of Slow CW Decrease for IEEE 802.11 WLAN. In *14th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC'03*, volume 2, pages 1717–1721, September 2003.

[237] A. Nicholson, Y. Chawathe, M. Chen, B. Noble, and D. Wetherall. Improved Access Point Selection. In *4th International Conference on Mobile Systems, Applications and Services, MOBISYS'06*, June 2006.

[238] B. O'Hara, P. Calhoun, and J. Kempf. Configuration and Provisioning for Wireless Access Points (CAPWAP) Problem Statement. RFC 3990 (Informational), February 2005.

[239] Inc. OPNET Technologies. *Wireless LAN Model User Guide*, 2004.

[240] E.A. Panaousis, P.A. Frangoudis, C.N. Ververidis, and G.C. Polyzos. Optimizing the channel load reporting process in IEEE 802.11k-enabled WLANs. In *16th IEEE Workshop on Local and Metropolitan Area Networks, LANMAN'08*, pages 37–42, September 2008.

[241] M.K. Panda, A. Kumar, and S.H. Srinivasan. Saturation throughput analysis of a system of interfering IEEE 802.11 WLANs. In *6th IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks, WoWMoM'05*, June 2005.

[242] Ioannis Papanikos and Michael Logothetis. A Study on Dynamic Load Balance for IEEE 802.11b Wireless LAN. In *8th International Conference on Advances in Communications and Control, COMCON'01*, June 2001.

[243] J. Paradells, M. Catalán, JL. Ferrer, M. Catalán-Cid, X. Sánchez, V. Beltrán, E. Garcia, C. Gómez, P. Plans, J. Rubio, D. Almodóvar, D. Rodellar, N. Subotic, D. Wenger, and I. Steiner. Red mallada asistida por UMTS/GPRS. In *VI Jornadas de Ingeniería Telemática, JITEL'07*, pages 17–24, 2007.

[244] J. Paradells, JL. Ferrer, M. Catalán, W. Torres, X. Sánchez, V. Beltrán, E. Garcia, C. Gómez, P. Plans, J. Rubio, D. Almodóvar, D. Rodellar, N. Subotic, D. Wenger, and I. Steiner. Design of a UMTS/GPRS Assisted Mesh Network (UAMN). In *17th Wireless World Research Forum Meeting, WWRF'06*, November 2006.

[245] M. Petrova and P. Mähönen. Cognitive Resource Manager: a cross-layer architecture for implementing cognitive radio networks. In F. Fittzek and M. Katz, editors, *Cognitive Wireless Networks*, chapter 20. Springer, 2007.

[246] N. Prasad and A. Prasad, editors. *WLAN Systems and Wireless IP for Next Generation Communications*. Artech House, co., London, UK, 2002.

[247] C. Prehofer and C. Bettstetter. Self-Organization in Communication Networks: Principles and Design Paradigms. *IEEE Communications Magazine*, 43(7):78–85, July 2005.

[248] John Proakis and Massoud Salehi. *Digital Communications*. McGraw Hill Higher Education, New York, NY, USA, 5th edition, January 2008.

[249] R.J. Punnoose, P.V. Nikitin, and D.D. Stancil. Efficient simulation of Ricean fading within a packet simulator. In *IEEE 52nd Vehicular Technology Conference, VTC'00-Fall*, volume 2, pages 764–767, 2000.

[250] M.B. Pursley and D.J. Taipale. Error Probabilities for Spread-Spectrum Packet Radio with Convolutional Codes and Viterbi Decoding. *IEEE Transactions on Communications*, 35(1):1–12, January 1987.

[251] D. Qiao and S. Choi. New 802.11h mechanisms can reduce power consumption. *IT Professional*, 8(2):43–48, March 2006.

[252] D. Qiao, S. Choi, and K.G. Shin. Goodput analysis and link adaptation for IEEE 802.11a wireless LANs. *IEEE transactions on Mobile Computing*, 1(4):278–292, December 2002.

[253] Daji Qiao, Sunghyun Choi, Amit Jain, and Kang G. Shin. Miser: an optimal low-energy transmission strategy for ieee 802.11a/h. In *9th annual international conference on Mobile computing and networking, MobiCom'03*, pages 161–175, New York, NY, USA, 2003. ACM Press.

[254] L. Qin and T. Kunz. Survey on Mobile Ad Hoc Network Routing Protocols and Cross-Layer Design. Technical Report SCE-04-14, Systems and Computer Engineering, Carleton University, August 2004.

[255] J. Rao and S. Biswas. Transmission power control for 802.11: a carrier-sense based NAV extension approach. In *IEEE Global Telecommunications Conference, GLOBECOM'05*, volume 6, pages 3439–3444, November 2005.

[256] Theodore S. Rappaport. *Wireless Communications Principles and Practices.* Prentice Hall PTR, Upper Saddle River, NJ, USA, 2nd edition, December 2001.

[257] P. Raptis, V. Vistas, K. Paparrizos, P. Chatzimisios, and A.C. Boucouvalas. Packet delay distribution of the IEEE 802.11 distributed coordination function. In *6th IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks. WoWMoM'05.*, pages 299–304, June 2005.

[258] J. Riihijärvi, M. Petrova, P. Mähönen, , and J. de Almeida Barbosa. Performance Evaluation of Automatic Channel Assignment Mechanism for IEEE 802.11 Based on Graph Colouring. In *The 17th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC'06.*, September 2006.

[259] J. Riihijärvi, M. Petrova, and P. Mähönen. Frequency Allocation for WLANs Using Graph Colouring Techniques. In *Second Annual Conference on Wireless On-demand Network Systems and Services, WONS'05*, pages 216–222, January 2005.

[260] R.C. Rodrigues, G.R. Mateus, and A.A.F Loureiro. On the design and capacity planning of a wireless local area network. In IEEE/IFIP, editor, *Network Operations and Management Symposium, NOMS-00*, pages 335–348, April 2000.

[261] M.T. Rose and K. McCloghrie. Structure and identification of management information for TCP/IP-based internets. RFC 1155 (Standard), May 1990.

[262] M.T. Rose and K. McCloghrie. Concise MIB definitions. RFC 1212 (Standard), March 1991.

[263] Oscar Santillana Ortega. Punts d'accés 802.11 intel·ligents basats en linux (II). Master's thesis, Escola Politècnica Superior de Castelldefels, EPSC - UPC, January 2005.

[264] M. Sedighizad, B. Seyfe, and K. Navaie. MR-BART: multi-rate available bandwidth estimation in real-time. In *Proceedings of the 3nd ACM workshop on Performance monitoring and measurement of heterogeneous wireless and wired networks, PM2HW2N'08*, pages 1–8, 2008.

[265] S-T. Sheu, Y-H. Lee, M-H. Chen, Y-C. Yu, and Y-C. Huang. PLFC: The packet length fuzzy controller to improve the performance of WLAN under the interference of microwave oven. In *Global Telecommunications Conference, GLOBECOM'00*, pages 1427–1431, November 2000.

[266] Y. Shi, D. Perkins, C. Elliott, and Y. Zhang. CAPWAP Protocol Base MIB. draft-ietf-capwap-base-mib-04, February 2009.

[267] Y. Shi, D. Perkins, C. Elliott, and Y. Zhang. CAPWAP Protocol Binding MIB for IEEE 802.11. draft-ietf-capwap-802dot11-mib-03, March 2009.

[268] B. Simsek, K. Wolter, and H. Coskun. Analysis of the qbss load element parameters of 802.11e for a priori estimation of service quality. *International Journal of Simulation: Systems, Science and Technology, Special Issue: Performance Engineering of Computer and Communication Systems*, 7(2), 2006.

[269] I. Siomina. *Radio Network Planning and Resource Optimization.* PhD thesis, Linköping University, 2007.

[270] B. Sklar. *Home Networking Basis: Transmission Environments and Wired/Wireless Protocols.* Prentice Hall, 2nd edition, January 2001.

[271] V. Srivastava and M. Motani. Cross-layer design: A survey and the road ahead. *IEEE Communications Magazine*, 43(12):112–119, 2005.

[272] Neng-Jian Tai, Jie Wu, Yan-Xin Gou, Yong-Min Wang, Cheng-Xi Dong, and Yan Tian. A New Approach to Analyze CCK Performance. In *International Conference on Wireless Communications, Networking and Mobile Computing, WiCom'07*, pages 1356–1360, September 2007.

[273] RW. Thomas, DH. Frlend, LA. Da Silva, and AB. MacKenzie. Cognitive Networks: Adaptation and Learning to Achieve End-to-End Performance Objectives. *IEEE Communications Magazine*, pages 51–57, December 2006.

[274] T.Nadeem and A. Agrawala. IEEE 802.11 DCF Enhancements for Noisy Environments. In *15th IEEE International Symposium On Personal, Indoor And Mobile Radio Communications, PIMRC'04, Barcelona*, September 2004.

[275] A. Tønnesen, A. Hafslund, and Ø . Kure. The Unik-OLSR Plugin Library. In *The OLSR Interop & Workshop*, August 2004.

[276] F.A. Tobagi and L. Kleinrock. Packet Switching in Radio Channels: Part II - the Hidden Terminal Problem in Carrier Sense Multiple Access and the Busy-Tone Solution. *IEEE Transactions on Communications*, pages 1417–1433, December 1975.

[277] Jean Tourrilhes. Wireless LAN resources for Linux. http://www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/Wireless.html, July 2007.

[278] S-C Tuan, J-C Chen, H-T Chou, and H-H Chou. Optimization of propagation models for the radio performance evaluation of wireless local area network. In IEEE, editor, *Antennas and Propagation Society International Symposium, 2003*, volume 2, pages 146–149, June 2003.

[279] S. Vasudevan, K. Papagiannaki, C. Diot, J. Kurose, and D. Towsley. Facilitating access point selection in IEEE 802.11 wireless networks. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement, IMC'05*, pages 26–26, 2005.

[280] Hector Velayos, Victor Aleo, and Gunnar Karlsson. Load balancing in overlapping wireless LAN cells. In *IEEE International Conference on Communications, 2004. ICC'04*, volume 7, pages 3833–3836, June 2004.

[281] Hector Velayos and Gunnar Karlsson. Multi-Rate Performance Index for Wireless LANs. In *15th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, 2004. PIMRC 2004*, volume 2, pages 1154–1157, September 2004.

[282] Hector Velayos and Gunnar Karlsson. Statistical analysis of the IEEE 802.11 MAC service time. In *Proceedings of the 19th International Teletraffic Congress, ITC19*, August 2005.

[283] Hector Velayos, Ignacio Más, and Gunnar Karlsson. Overload protection for IEEE 802.11 cells. In *14th IEEE International Workshop on Quality of Service, IWQoS'06*, June 2006.

[284] Hector Velayos-Muñoz. *Autonomic Wireless Networking.* PhD thesis, KTH Royal Institute of Technology, 2005.

[285] V.M. Vishnevsky and A.I. Lyakhov. 802.11 LANs: Saturation Throughput in Presence of Noise. In *Proc. of IFIP Networking*, May 2002.

[286] V. Vitsas and A. Boucouvalas. Performance Analysis of the Advanced Infrared (AIr) CSMA/CA MAC Protocol for Wireless LANs. *ACM Wireless Networks*, 9:495–507, September 2003.

[287] S-Y Wang, C.L. Chou, C.H. Huang, C.C. Hwang, Z.M. Yang, C.C. Chiou, and C.C. Lin. The Design and Implementation of the NCTUns 1.0 Network Simulator. *Computer Networks*, 42(2):175–197, June 2003.

[288] Shie-Yuan Wang and Yi-Bing Lin. Nctuns network simulation and emulation for wireless resource management. *Wireless Communications and Mobile Computing*, 5(8):899–916, December 2005.

[289] Y. Wang, L.G. Cuthbert, and J. Bigham. Intelligent Radio Resource Management for IEEE 802.11 WLAN. In *IEEE Wireless Communications and Networking Conference, WCNC'04*, volume 3, pages 1365–1370, March 2004.

[290] S-Y Wang et. al. The NCTUns network simulator and emulator v5.0. http://nsl10.csie.nctu.edu.tw/, 2008.

[291] P. Wertz, M. Sauter, G. Wölfle, R. Hoppe, and FM. Landstorfer. Automatic optimization algorithms for the planning of wireless local area networks. In *IEEE 60th Vehicular Technology Conference, VTC'04-Fall*, September 2004.

[292] Roger M. Whitaker and Steve Hurley. Evolution of planning for wireless communication systems. In *Proceedings of the 36th Annual Hawaii International Conference (HICSS'03), Big Island, HI.*, pages 296–305. IEEE Computer Society, 2003.

[293] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *Proceedings of the 2nd IEEE workshop on Embedded Networked Sensors, EmNets'05*, pages 45–52. IEEE Computer Society, 2005.

[294] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma. Performance of Reliable Transport Protocol over IEEE 802.11 Wireless LAN: Analysis and Enhancement. In *The 21st Conference on Computer Communications, INFOCOM'02*, volume 2, pages 599–607, June 2002.

[295] M. Daoud Yacoub. Fading distributions and co-channel interference in wireless systems. *Antennas and Propagation Magazine, IEEE*, 42(1):150–160, February 2000.

[296] L. Yang, P. Zerfos, and E. Sadot. Architecture Taxonomy for Control and Provisioning of Wireless Access Points (CAPWAP). RFC 4118 (Informational), June 2005.

[297] L.H. Yen, T.T. Yeh, and K.H. Chi. Load Balancing in IEEE 802.11 Networks. *IEEE Internet Computing*, 13(1):56–64, 2009.

[298] Li-Hsing Yen and Tse-Tsung Yeh. SNMP-Based Approach to Load Distribution in IEEE 802.11 Networks. In *IEEE 63rd Vehicular Technology Conference, VTC'06-Spring*, volume 3, pages 1196–1200, May 2006.

[299] J. Zhu and S. Roy. 802.11 Mesh Networks with Two Radio Access Points. In *IEEE International Conference on Communications, 2005. ICC'05*, volume 5, pages 3609–3615, May 2005.

[300] Rodger E. Ziemer and Roger W. Peterson. *Introduction to Digital Communication.* Prentice Hall PTR, Upper Saddle River, NJ, USA, 2nd edition, August 2000.

[301] E. Ziouva and T. Antonakopoulos. The IEEE 802.11 Distributed Coordination Function in Small-Scale Ad-Hoc Wireless LANs. *International Journal of Wireless Information Networks*, 10(1):1–15, January 2003.

Abbreviations and Acronyms

AAA	Authentication, Authorization and Accounting
AAC	Available Admission Capacity
ACL	asynchronous connectionless
ACK	Acknowledgment
ADC	Analog-to-Digital Converter
AES	Advanced Encryption Standard
AFH	adaptive frequency-hopping
AGC	Automatic Gain Control
AIFS	Arbitration Interframe Space
AODV	Ad hoc On-Demand Distance Vector
AP	Access Point
API	Application Programming Interface
APSD	Automatic Power Save Delivery
ARP	Address Resolution Protocol
ASN.1	Abstract Syntax Notation One
AWGN	Additive White Gaussian Noise
AWMA	Alternating Wireless Medium Access
BER	Bit Error Ratio
BO	Backoff
BPSK	Binary Phase-Shift Keying
BS	Base Station
BSS	Basic Service Set
BSSID	BSS Identification
CAPWAP	Control and Provisioning of Wireless Access Points
CBR	constant bitrate
CCA	Clear Channel Assessment
CCK	Complementary Code Keying
CCMP	Counter mode with Cipher-block chaining Message authentication code Protocol
CL	Cross-Layer
CMIP	Common Management Information Protocol
CR	Cognitive Radio
CRC	Cyclic Redundancy Check
CS	Carrier Sense

CSMA/CA	Carrier Sense Multiple Access with Collision Avoidance
CTS	Clear to Send
CW	Contention Window
DARPA	Defense Advanced Research Projects Agency
DBPSK	Differential BPSK
DCF	Distributed Coordination Function
DFS	Dynamic Frequency Selection
DIFS	DCF Interframe Space
DIMACS	Discrete Mathematics and Theoretical Computer Science
DLS	Direct Link Setup
DQPSK	Differential QPSK
DS	Distribution System
DSA	Dynamic Spectrum Allocation
DSATUR	Degree of Saturation
DSSS	Direct Sequence Spread Spectrum
DTLS	Datagram Transport Layer Security
EDCA	Enhanced DCF Channel Access
EIRP	Equivalent Isotropically Radiated Power
Elwis	Elena's WLAN incredible simulator
ERO	European Radiocommunications Office
ERP	Extended Rate PHY
ESS	Extended Service Set
ESSID	ESS Identification
ETSI	European Telecommunications Standards Institute
FAP	Frequency Assignment Problem
FCC	Federal Communications Commission
FCS	Frame Check Sequence
FHSS	Frequency-Hopping Spread Spectrum
FTP	File Transfer Protocol
GCC	GNU Compiler Collection
GL	Gross Load
glibc	GNU C Library
GNU	GNU's not Unix
GPL	GNU General Public License
GPRS	General Packet Radio Service
GUI	graphical user interface
HCCA	HCF Controlled Channel Access
HCF	Hybrid Coordination Function
HiperLAN	HIgh PErformance Radio LAN
HR-DSSS	High Rate DSSS
HTTP	HyperText Transfer Protocol

I^{2}C	Inter-Integrated Circuit
IAPP	Inter Access Point Protocol
IBSS	Independent Basic Service Set
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ILP	Integer Linear Programming
ioctl	Input/output control
IR	Infra Red
ISI	intersymbol interference
ISM	Industrial Scientific Medical
ITS	Intelligent Transportation Systems
ITU	International Telecommunication Union
LAN	Local Area Network
LLC	Logical Link Control
LTE	Long Term Evolution
LWAPP	Light Weight Access Point Protocol
MAC	Medium Access Control
MAD	Medium Access Delay
MbD	Management by Delegation
MIB	Management Information Base
MIMO	multiple-input multiple-output
MIPS	Microprocessor without Interlocked Pipeline Stages
MIT	Massachusetts Institute of Technology
MLME	MAC Sublayer Management Entity
MPDU	MAC Protocol Data Unit
MPR	Multi-Point Relay
MQOK	M-ary Quaternary Orthogonal Keying Modulation
MSDU	MAC Service Data Unit
NAM	Network Animator
NAV	Network Allocation Vector
NIC	Network Interface Card
ns-2	Network Simulator 2
NSF	National Science Foundation
NSTA	number of associated stations
NTIA	National Telecommunications and Information Administration
OFDM	Orthogonal Frequency-Division Multiplexing
OH	overhead
OID	object identifier
OLSR	Optimized Link State Routing Protocol
OS	Operating System
OSPF	Open Shortest Path First

PAM	Pulse Amplitude Modulation
PBCC	Packet Binary Convolutional Coding
PC	Personal Computer
PCF	Point Coordination Function
PCI	Peripheral Component Interconnect
PCMCIA	Personal Computer Memory Card International Association
PER	Packet Error Ratio
PHY	Physical Layer
PIFS	PCF Interframe Space
PL	Packet Loss
PLCP	Physical Layer Convergence Procedure
PMD	Physical Medium Dependent
PPDU	PLCP Protocol Data Unit
PSD	power spectral density
PSM	Power Saving Mode
PTA	Packet Traffic Arbitration
QAM	Quadrature Amplitude Modulation
QAP	QoS enabled AP
QoE	Quality of Experience
QoS	Quality of Service
QPSK	Quadrature Phase-Shift Keying
RF	radio frequency
RFC	Request for Comments
RPI	Receive Power Indication
RRM	radio resource management
RSN	Robust Security Network
RSNA	Robust Security Network Association
RSSI	Received Signal Strength Indicator
RTS	Request to Send
RTT	round trip time
SCO	synchronous connection oriented
SDR	Software Defined Radio
SER	Symbol Error Ratio
SFD	Start Frame Delimiter
SIFS	Short Interframe Space
SINR	Signal to Interference and Noise Ratio
SMI	Structure of Management Information
SMX	spatial multiplexing
SNMP	Simple Network Management Protocol
SNR	Signal-to-Noise Ratio
SOHO	Small Office/Home Office
STA	Station
SYNC	Synchronization

TBRPF	Topology Dissemination Based on Reverse-Path Forwarding
Tcl	Tool Command Language
TCP	Transmission Control Protocol
TG	Task Group (within the IEEE 802.11 WG)
TKIP	Temporal Key Integrity Protocol
TMT	theoretical maximum throughput
TPC	Transmit Power Control
TXOP	Transmission Opportunity
UAMN	UMTS-Assisted Mesh Network
UCLA	University of California at Los Angeles
uClibc	μC Library
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunications System
UNII	Unlicensed National Information Infrastructure
USC/ISI	University of Southern California/Information Science Institute
UWB	Ultra Wide Band
WAN	Wide Area Network
WAVE	Wireless Access in the Vehicular Environment
WDS	Wireless Distribution System
WEP	Wired Equivalent Privacy
WEXT	Wireless Extensions
WG	Working Group (within IEEE 802 LAN/MAN Standards Committee)
Wi-Fi	Wireless Fidelity
WLAN	Wireless LAN
WMN	Wireless Mesh Network
WPA	Wi-Fi Protected Access
WPAN	Wireless Personal Area Network
YANS	Yet Another Network Simulator