

1 A cognitive radio world

1.1 Introduction

Cognitive radio is a topic of great interest and holds much promise as a technology that will play a strong role in communication systems of the future. This book focuses on the essential elements of cognitive radio technology and regulation. This is a challenging task in that cognitive radio is still very much an emerging technology. There is much debate over its exact definition, its potential role in communication systems, whether cognitive radios should in fact be permitted in the first place and if yes, what the regulatory policies should be. However, while acknowledging the flux in this field, the book aims to identify the core concepts that will remain central to the field irrespective of how precisely it develops. The aim of this first chapter is to briefly define cognitive radio and to then focus on the all important question of why cognitive radios are needed. This chapter therefore motivates all that is to come in the book.

1.2 Brief history and definition

The term cognitive radio was coined by Mitola in an article he wrote with Maguire in 1999 [1]. In that article, Mitola and Maguire describe a cognitive radio as a radio that understands the context in which it finds itself and as a result can tailor the communication process in line with that understanding. Since the coining of the phrase, the term cognitive radio has grown and expanded and has tended to be used in very many ways. With this in mind it is perhaps useful to spend some time defining cognitive radio and place that definition in the context of current radio technologies.

In very simple terms a cognitive radio is a *very smart radio*. Radios have been getting smarter in the last few decades. Current communication

2 A COGNITIVE RADIO WORLD

systems use radios that can adapt their behaviour in many ways. For example 3G communication devices have the ability to dynamically alter their power output in order to ensure power imbalances, which negatively affect communication, do not arise between different users. Mobile phones can cleverly process the incoming signals they receive in order to mitigate against the various different distortion effects that the signal experiences. WiMAX¹ systems can adapt the characteristics of signals they transmit in order to maintain good throughput and link stability. All this type of functionality is invisible to the user, but the fact remains that the communication systems we use today are able to adapt and change their behaviour in a variety of ways to maintain connectivity in the face of varying conditions and circumstances.

In all of the above examples, the adaptations that occur are well defined and can be anticipated. And the adaptations are triggered by straightforward and well-understood conditions. Let us look at the WiMAX example to see what this means. Modulation is the process by which data is placed on the radio waves for transmission. The order of the modulation scheme gives an indication of how compactly the data is modulated onto the waves. Higher-order schemes get more data through but need good signal conditions to work. Lower-order modulation schemes get less data through but need less good signal conditions. A received signal is typically very good near the basestation and therefore a high-order modulation scheme can be used. However, in areas close to the edge of the range of the WiMAX basestation, the received signal is much poorer. So, the system steps down to a lower-order modulation scheme to maintain the connection quality and link stability. Hence the modulation scheme changes with distance from the basestation as indicated in Figure 1.1. The modulation scheme is known as an adaptive modulation scheme and there is a simple mapping between the quality of the signal at the receiver and the modulation scheme used.

A cognitive radio takes this type of adaptive behaviour and goes much further. By this we mainly mean two things. Firstly the level of adaptivity

1 WiMAX is a wireless digital communications system, also known as IEEE 802.16, that is intended for wireless metropolitan area networks.




Fig. 1.1. Adaptive modulation in action for a WiMAX system.

is greatly increased and applies to as wide a range of operating parameters as possible, e.g. frequency of operation, power, modulation scheme, antenna beam pattern, battery usage, processor usage, etc. And secondly the adaptation itself can happen in both planned and unplanned ways. The latter can be made possible through the radio recognising patterns of behaviour, learning from reoccurring situations and past experiences and using mechanisms for anticipating future events. With this in mind, a cognitive radio can be defined.

A cognitive radio is a device which has four broad inputs, namely, an understanding of the environment in which it operates, an understanding of the communication requirements of the user(s), an understanding of the regulatory policies which apply to it and an understanding of its own capabilities. In other words a cognitive radio is fully *aware* of the context in which it is operating. A cognitive radio processes the inputs it receives and makes *autonomous decisions* on how to configure itself for the communication tasks at hand. In deciding how to configure itself, the radio

4 A COGNITIVE RADIO WORLD

attempts to match actions to requirements while at the same time being cognisant of whatever constraints or conflicts (physical, regulatory, etc.) may exist. A cognitive radio has the ability to *learn* from its actions and for this learning to feed into any future reactions it may have. A cognitive radio is made from software and hardware components that can facilitate the wide variety of different configurations it needs to communicate.

This definition, while perhaps not suitable for an elevator pitch, does capture the essential ingredients of a cognitive radio.² *How a cognitive radio gets the input it needs, processes the information, decides on how to configure itself, puts its decisions into action and deals with learning is very much the focus of the book.* However, before addressing these issues, the question that needs to be addressed first is, ‘Why do we need cognitive radios?’ A cognitive radio is undoubtedly more complex than any existing radio and the need for the extra complexity must be justified.

A succinct way of answering this question is that we will be able to do ‘new things that we currently cannot do’ and we will be able to do ‘old things better’ with a cognitive radio. The new things that we will be able to do are very much connected with new spectrum management regimes. The old things that we will be able to do better lie mainly in the domain of autonomous organisation and management of increasingly complicated communication systems. And in general it is fair to say that cognitive radios will make communication possible in ever more stressful and challenging circumstances.

1.3 New spectrum regimes

We start by looking at the role of cognitive radio in spectrum management. We do this for two reasons. The first is that historically the motivating applications for cognitive radio have been presented in the context of spectrum management and in particular in relation to *dynamic spectrum access*, a concept which will become clear later in this chapter. The second reason for starting with spectrum management is that many of the wider

2 Currently efforts are being made to more tightly define the term. So for example the IEEE 1900.1 Working Group on Terminology and Concepts for Next Generation Radio Systems and Spectrum Management has developed a set of standard definitions.

applications for cognitive radio make use of the dynamic techniques that are fundamental to dynamic spectrum management concepts.

In very general terms, spectrum management involves the process of organising how the spectrum is used and by whom. The key purpose of spectrum management is to maximise the value that society gains from the radio spectrum by allowing as many efficient users as possible while ensuring that the interference between different users remains manageable [2]. New dynamic spectrum management regimes are on the horizon as we move away from the current static approaches and cognitive radio has an enabling role to play.

1.3.1 Current regimes

Historically, the approach adopted by spectrum managers around the world to managing the radio spectrum has been highly prescriptive. Typically, regulators decide on the use of a particular range of frequencies, or frequency band, as well as specifying what services should be delivered in the band, which technologies are permitted in the delivery of the services and who gets to deliver and perhaps use the services. This is referred to as the *administrative approach* to spectrum management. The term *Command and Control* is also often used to describe this approach, or at least to delineate traditional practices which do not necessarily heed market demands, in spectrum management.

The entire radio spectrum is divided into blocks, or bands, of frequencies established for a particular type of service by the process of frequency *allocation*. Frequency allocation is performed on an international and national basis. Broadly speaking, international bodies tend to set out high-level guidance, to which national bodies adhere in setting more detailed policy. At the highest level of management sits the International Telecommunication Union (ITU), a specialised agency of the United Nations. The ITU's International Radio Regulations allocate the spectrum from 9 kHz to over 275 GHz to a range of different uses. In some cases, there are multi-national bodies coordinating across a region. A good example of this is the case of the European Union (EU). At national level typically a national regulatory authority (NRA) manages the day-to-day

6 A COGNITIVE RADIO WORLD

use of the spectrum, in line with ITU guidelines and for example, in the case of the EU member states, in line with EU policy. Within the broad international frequency allocations, the national regulatory authority can make more specific channel plans. For example, allocations made to the land mobile service can be divided into allotments for business users, public safety users and cellular users, with each group allotted a portion of the band in which to operate. The upshot of this approach is that blocks of the spectrum are set aside for specific activities in each country and the regulator knows what to expect.

The tight control over the *use* of the spectrum is attractive to regulators because with this approach interference can be managed. All communication systems cause a degree of interference. A full description of how interference arises can be found in Chapter 3 but for the moment we can simply think of interference as meaning a distortion of the transmitted waves that can actually prevent the receiver from being able to correctly decipher the incoming signal. In broad terms interference arises when services in neighbouring frequency bands or services in the same frequency bands but in different geographical areas interact with each other in an undesirable way. The administrative approach makes it easier for the regulator to ensure that excessive interference does not occur because the regulator is able to carefully model the interaction between services in neighbouring bands and in different geographical areas and tailor the licence conditions appropriately. This tailoring, for example, manifests itself in the specification of guard bands between services or conditions on maximum power transmission levels. Guard bands are bands of frequencies which are deliberately left free to ensure neighbouring services do not *spill* over into each other. Figure 1.2 captures this concept and shows tightly proscribed neighbouring services, and the associated regulated guard bands, over a range of frequencies.³ In Figure 1.2 two neighbouring services are defined as neighbours from a frequency perspective. Two neighbouring services can also be defined in a spatial context and it is in this context that the

3 This diagram is inspired by a diagram in an Ofcom report which can be downloaded from <http://www.ofcom.org.uk/consult/condocs/sur/spectrum/>.




Fig. 1.2. Tightly administered allocations of frequency bands.

power restrictions mentioned above come into play. Chapter 3 will deal with this in more detail.

So, once allocation has taken place, *assignment* is the next stage of the spectrum management process. Assignment, which happens at national level, refers to the final subdivision of the spectrum in which the spectrum is actually assigned to a specific party for use. A wide range of mechanisms exist for making assignments. In cases where there is limited demand for spectrum, something as simple as a first come first served can be used. On the other hand, in cases of high demand, more complex mechanisms are used. For example, in the case of the 3G bands, regulators around the world used beauty contests⁴ and auctions to assign the spectrum in the specific bands set aside for these services. Not all spectrum costs money. In the 3G bands just mentioned, large sums of money were handed over for the assignments. On the other hand, TV broadcasters and military users do not pay for spectrum in many countries. The essential point here is not to provide a comprehensive understanding of current spectrum management rules and regulations but rather to underline the fact that centralised decisions, both technical and economic, regarding usage of the spectrum are a key part of the administrative approach to

⁴ A beauty contest is a merit-based comparative evaluation approach in which interested parties make a submission and are judged and ranked.

spectrum management. Cave *et al.* [2] provide a very good overview of all matters spectrum management.

While the administrative approach has its attractions for regulators and a degree of certainty for incumbents and users, it has many disadvantages. The spectrum planning involved in allocating frequency bands to certain uses and then regulating what equipment can be used to deliver the services can be a slow process and not capable of keeping up with new innovations and technologies. Potential users of spectrum can make proposals for allocations, for example for new communication technologies, but without the allocation being made, matters cannot progress further. So, for example, operators who received licences to provide 3G services are not able to use their allocated spectrum for other services while the 3G market builds. Nor are they easily able to divert 2G spectrum to growing 3G service demands (though this situation is currently being addressed). The frequency allocation process itself forces regulators to ‘pick winners’. This can be a very difficult task and it is not sensible or efficient to put the regulator in such a position.⁵ There tend to be limited incentives for those who have spectrum to use it efficiently while spectrum is seen as a scarce resource to those in search of it.

However, while on the one hand spectrum is considered scarce, on the other hand there is another story to tell. Figure 1.3 shows a set of spectrum occupancy measurements taken over a particular range of frequencies in a particular city centre location in Dublin City in April 2007. The measurements involved were taken by the Shared Spectrum Company using dedicated equipment and cover a time span which begins at 6 pm on one evening and continues for 40 hours. This image is typical of very many spectrum occupancy measurements⁶ in that it exhibits lots of unused spectrum. And in fact figures such as spectrum occupancy levels of as little as 10% are often suggested as giving the real picture of spectrum usage.

5 An example of where this has gone wrong is the allocation of spectrum to the ERMES paging system or the TFTS in-flight phone system in Europe. These allocations have resulted in spectrum being unused for over a decade.

6 Shared Spectrum Company have a wide range of such measurements taken in different cities, available on their website <http://www.sharedspectrum.com/>.




Fig. 1.3. Shared Spectrum Company measurements, Dublin, April 2007.

The middle plot of Figure 1.3 is the one that deserves attention. Spectrum occupancy between 1400 and 1520 MHz is shown.⁷ As can be seen there are many frequencies in which there are no transmitters present. The temporal nature of the transmissions can also be clearly seen. In particular the transmissions that occur in the 1500 MHz to 1520 MHz range exhibit strong temporal behaviour with significant unoccupied spectrum during two time intervals which in fact correspond to the two nights over which the measurements were taken. The key point to be made is that, despite the fact that all the frequencies in the above bands have been allocated to particular services, all the bands are not fully occupied.

The recognition of the many short-comings in administrative approaches to spectrum management has precipitated a search for new

⁷ Note that spectrum occupancy in Figure 1.3 is the spectrum occupancy seen by a radio at a given location and another radio elsewhere may get a somewhat different view.

spectrum management regimes and mechanisms. There are many different innovations on the horizon and there is much discussion and debate and much research as to which approaches are best. What is of interest in this book are the new spectrum management regimes which are relevant to cognitive radio. We begin by looking at *dynamic spectrum access* regimes.

1.3.2 Dynamic spectrum access regimes

Measurements such as those in Figure 1.3 and many others have led to the suggestion that it should be possible to access spectrum dynamically. What this means is that radios, rather than being given a static range of frequencies on which to operate, could instead be allowed to use whatever unused spectrum or *white space* they find free. This is essentially what is meant by *dynamic spectrum access*. The dynamic approach has the potential to make use of what are otherwise wasted resources. The opening up of spectrum in this manner also has the potential to reduce barriers to entry for new ventures as this technique means more spectrum is available for use.

In the first instance the vision for dynamic spectrum access is based on one in which new users ‘fill’ unoccupied spectrum around existing users. As can be seen from Figure 1.3 some frequencies will be unoccupied for long durations while others are only unoccupied overnight. Others still, not obvious from Figure 1.3, will be much shorter lived and perhaps last minutes and seconds. Needless to say different types of applications may suit longer-term or shorter-term gaps.

In the main, the dynamic spectrum access vision has centred around the concept of spectrum sharing between licensed and unlicensed users. In this case, the holders of licences to the spectrum are known as *primary users*. These licensed users have priority access to the spectrum for which they have a licence and *secondary users* can use the spectrum when the primary user does not need it. The secondary users cease transmitting on the return of the primary user. For this to work, the secondary user needs to be able to detect the white space, configure itself to transmit in that white space, detect the return of the primary user and then