CHAPTER

When radio meets software

1

Alexander M. Wyglinski[1], Maziar Nekovee[2], and Y. Thomas Hou[3]
[1]Worcester Polytechnic Institute, United States
[2]BT Research and University College London, United Kingdom
[3]Virginia Polytechnic Institute and State University, United States

1.1 INTRODUCTION

Data communication networks are a vital component of any modern society. They are used extensively in numerous applications, including financial transactions, social interactions, education, national security, and commerce. In particular, both wired and wireless devices are capable of performing a plethora of advanced functions that support a range of services, such as voice telephony, web browsing, streaming multimedia, and data transfer. With the rapid evolution of microelectronics, wireless transceivers are becoming more versatile, powerful, and portable. This has enabled the development of *software-defined radio* (SDR) technology, where the radio transceivers perform the baseband processing entirely in software: modulation/demodulation, error correction coding, and compression.

Since its introduction in 1991, SDR has been defined as a radio platform of which the functionality is at least partially controlled or implemented in software. Consequently, any waveform defined in the memory of the SDR platform can be employed on any frequency [1]. Although initially constrained by the conversion process between the analog and digital signaling domains, the emergence of cheap high-speed digital-to-analog converters (DACs) and analog-to-digital converters (ADCs) has brought the ideal SDR concept of an entirely software communication system implementation (including radio frequency functionality) closer to a reality.

Wireless devices that can be described as SDR have in fact been around for several decades. They were initially employed in military applications before finding applications in the commercial sector. Military programs such as SPEAKeasy sought to enable communication and interoperability between several military standards [2]. Although ambitious, the SPEAKeasy project did produce a functional prototype, even though the design choices involved in programming waveforms using low-level assembly language meant that the software was not compatible

Doi: 10.1016/B978-0-12-374715-0.00001-0

with newer processors. Furthermore, in terms of portability, the Phase I prototype of SPEAKeasy was large enough to fit into the back of a truck [3].

One of the first significant commercial introductions of SDR platforms was the Vanu Anywave™ software radio base station, which incorporated multiple cellular access standards into a simple SDR implementation. Since the cellular standards are based in software, they can be changed "on the fly" to adapt to different user needs of each cell, rather than replacing the radio frequency (RF) hardware, which can be a prohibitively expensive upgrade. Furthermore, new standards can be uploaded to the SDR platform for immediate deployment in a cellular region [4]. To increase the effectiveness and improve the aging process of an SDR platform, most developers seek to use portable code for their software, reusable components that can work under different waveform configurations, and generic hardware that can be easily upgraded [5].

Given the ease and speed of programming baseband operations in an SDR platform, this technology is considered to be a prime candidate for numerous advanced networking applications and architectures that were unrealizable only several years ago. An SDR platform that can rapidly reconfigure operating parameters based on changing requirements and conditions and through a process of cognition is known as *cognitive radio* [6]. The term *cognitive radio* (CR) was first defined by Joseph Mitola III [7]. According to Mitola, CR technology is the "intersection of personal wireless technology and computational intelligence," where CR is defined as "a really smart radio that would be self-aware, RF-aware, user-aware, and that would include language technology and machine vision along with a lot of high-fidelity knowledge of the radio environment" [7]. Cognitive radio clearly goes hand in hand with SDR; together, they can achieve functionality considered impossible only a decade ago. Consequently, before continuing any further with respect to CR, we first provide an overview of SDR technology.

1.2 SOFTWARE-DEFINED RADIO

1.2.1 What Is Software-Defined Radio?

Before describing what SDR does, it is useful to review the design of a conventional digital radio. Figure 1.1 shows a block diagram of a generic digital radio [8], which consists of five sections:

- The antenna section, which receives (or transmits) information encoded in radio waves.
- The RF front-end section, which is responsible for transmitting/receiving radio frequency signals from the antenna and converting them to an intermediate frequency (IF).
- The ADC/DAC section, which performs analog-to-digital/digital-to-analog conversion.

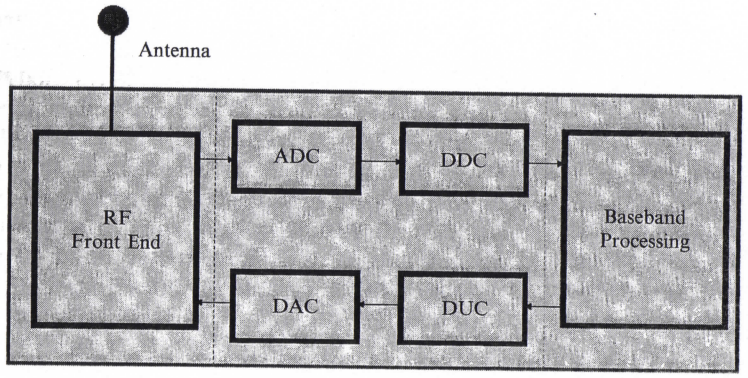

FIGURE 1.1

Schematic block diagram of a digital radio [8].

- The digital up-conversion (DUC) and digital down-conversion (DDC) blocks, which essentially perform modulations of the signal on the transmitting path and demodulation of the signal on the receiving path.
- The baseband section, which performs operations such as connection setup, equalization, frequency hopping, coding/decoding, and correlation, while also implementing the link layer protocol.

The DDC/DUC and baseband processing operations require large computing power, and in a conventional digital radio are implemented in dedicated hardware. In programmable digital radio (PDR) systems baseband operations and link layer protocols are implemented in software while the DDC/DUC functionality is performed using application-specific integrated circuits (ASICs).

Software-defined radio refers to technologies wherein these functionalities are performed by software modules running on field programmable gate arrays (FPGAs), digital signal processors (DSP), general-purpose processors (GPP), or a combination thereof. This enables programmability of both DDC/DUC and baseband processing blocks. Hence, operation characteristics of the radio, such as coding, modulation type, and frequency band, can be changed at will, simply by loading a new software. Also multiple radio devices using different modulations can be replaced by a single radio device that can perform the same task.

If the AD/DA conversion can be pushed further into the RF block, the programmability can be extended to the RF front end and an ideal *software* radio can be implemented. However, there are a number of challenges in the transition from hardware radio to software (-defined) radio. First, transition from hardware to software processing results in a substantial increase in computation, which in turn results in increased power consumption. This reduces battery life and is one of the key reasons why software-defined radios have not been deployed yet in end-user

devices, but rather in base stations and access points, which can take advantage of external power resources.

Second, the question where the AD/DA conversion can be performed determines what radio functions can be done in software and hence how reconfigurable a radio can be made. The ultimate goal for software radio is to move the AD/DA conversion as close as possible to the antenna so that all signal processing can be done digitally. However, two technical limitations make it currently infeasible to the AD/DA conversion at the antenna. First, digitization of the RF signal requires the incoming signal to be sampled at least at a rate that is determined by the Nyquist frequency. Additionally, the higher the data rate of the signal, the higher the resolution required to capture the information. Taken together, this means that high-bandwidth, high-frequency RF transmissions require very high sampling rates.

The ability to support very high sampling rates, which is especially critical with the use of high-frequency signals in the gigahertz range, limits the range of what can be digitized. To give an example, the typical channels used by an 802.11 WiFi device are 20 MHz wide. To assure that the full 20 MHz is presented to the modem without distortion, it is not unusual for ADC to digitize 40 MHz or so of signal bandwidth. To capture 40 MHz of analog signal bandwidth set by the IF filters without aliasing artifacts, the ADC will probably sample the signal at a rate above 80 million samples per second (Msps). Indeed, it is only recently that sufficiently fast DSPs and wideband AD/DA chipsets have become available at affordable cost to make it feasible to contemplate AD conversions of the IF rather than the baseband signal.

SDR is currently used to build radios that support multiple interface technologies (e.g., CDMA, GSM, and WiFi) with a single modem by reconfiguring it in software. However, SDR modems are expensive, since they typically entail programmable devices like FPGAs, as opposed to the mass-produced, single-purpose ASICs used in most consumer devices today (and are key enablers for low-cost handsets). Even today's multimode devices tend to just have multiple ASICs (or multiple cores on a single ASIC). SDR is currently used mostly in military applications, where cost is less of a constraint. SDR is also a modem technology and it ignores RF design issues. In particular, the RF design of a wireless device is typically closely coupled with the underlying access technology and modem design. For example, different air interface technologies have different spectral mask requirements and different degrees of vulnerability to cochannel interference and strong adjacent channel power. A device that must work over a wide bandwidth or over a wide range of RF signal scenarios (i.e., what other devices are operating in the nearby spectrum neighborhood) will be more complex and expensive than a single-purpose device.

1.2.2 Evolution of Software-Defined Radio

Two decades ago most radios had no software at all, and those that had it didn't do much with it. In a remarkably visionary article published in 1993 [2], Joseph Mitola III envisioned a very different kind of radio: A mostly digital radio that could

be reconfigured in fundamental ways just by changing the software code running on it. He dubbed this *software-defined radio*.

A few years later Mitola's vision started to become reality. In the mid-1990s military radio systems were invented in which software controlled most of the signal processing digitally, enabling one set of hardware to work on many different frequencies and communication protocols. The first (known) example of this type of radio was the U.S. military's SPEAKeasy I and SPEAKeasy II radios, which allowed units from different branches of armed forces to communicate for the first time. However, the technology was costly and the first design took up racks that had to be carried around in a large vehicle. SPEAKeasy II was a much more compact radio, the size of two stacked pizza boxes, and was the first SDR with sufficient DSP resources to handle many different kinds of waveforms [9]. SPEAKeasy II subsequently made its way into the U.S. Navy's digital modulator radio (DMR) with many waveforms and modes, able to be remotely controlled with an Ethernet interface. These SPEAKeasy II and DMR products evolved not only to define these radio waveform features in software, but also to develop an appropriate software architecture to enable porting the software to an arbitrary hardware platform, thus achieving independence of the waveform software specification and design from the underlying hardware [9].

In the late 1990s SDR started to spread from the military domain to the commercial sector, with the pace of penetration into this market considerably accelerating in the new millennium. Cellular networks were considered as the most obvious and potentially most lucrative market that SDR could penetrate. The benefits it could bring to this industry included a general-purpose and therefore more economic hardware platform, future-proofing and easier bug fixes through software upgrades, and increased functionality and interoperability through the ability to support multiple standards [10].

Companies such as Vanu, AirSpan, and Etherstack currently offer SDR products for cellular base stations. Vanu Inc., a U.S.-based company, has been focusing on the commercial development of SDR business since 1998. It received a lot of attention in 2005 with its Anywave™ GSM base station, which became the first SDR product to receive approval under the newly established software radio regulation. The Anywave base station runs on a general-purpose processing platform and provides a software implementation of the BTS (base transceiver station), BSC (base station controller), and TRAU (transcoder and rate adaptation unit) modules of the BSS (base sation subsystem). It supports GSM and can be upgraded to GPRS and Edge. The product was first deployed in rural Texas by Mid Tex Cellular in a trial, where Vanu base station showed successfully how it could concurrently run a time division multiple access (TDMA) and a GSM network, as well as remotely upgrade and fix bugs on the base station via an Internet link.

Following this successful trial other operators, such as AT&T and Nextel, expressed interest in the Anywave base station. In 2001 the 3GNewsroom was reporting SDR base stations as the key solution to the 3G rollout problem. The ability of SDR base stations to reconfigure on the fly and support multiple protocols was

thought to be the safest option for rolling out 3G. In reality, SDR didn't play the key role that was anticipated. However, a closer look at the operator's infrastructure shows that programmable devices have become a key component of current 3G base stations. In March 2005 Airspan released the first commercially available SDR-based IEEE 802.16 base station. The AS.MAX base station uses picoarrays™ and a reference software implementation of the IEEE 802.16d standard. The picoarray is a reconfigurable platform that is 10 times faster in processing power than today's DSPs. The AS.MAX base station promises to be upgradeable to the next generation mobile 802.16e standard and so has the potential to offer a future-proof route to operators looking to rolling out WiMAX services.

In addition to the preceding proprietary SDR platforms developed for the military and commercial sectors, there has also been significant progress in the SDR development in the open-source research and university communities. GNU Radio is an open-source architecture designed to run on general-purpose computers. It is essentially a collection of DSP components and supports RF interface to the universal software radio peripheral (USRP), an up- and down-convertor board coupled with ADC and DAC capabilities, which can be coupled to a daughter RF board. GNU Radio has been extensively used as an entry-level SDR within the research community. Some major SDR platforms developed in the university and research communities will be described in detail in this book.

As mentioned, due to its high demand on computation and processing, SDR technology has worked only in devices that have less constraint in size and power consumption, such as base stations and moving vehicles. But there is an increasing demand for SDR to enter portable and handheld devices in the future. The main issue with introducing SDR into portable devices has been that it requires the use of programmable platforms, which are generally power hungry and hence lead to reduced battery life and large devices. However, SDR provides the ability to support multiple waveforms on a single device, and so ultimately could give an end user increased choice of services if incorporated into a portable device, such as a handset. SDR could also assist seamless roaming at the national and international levels. However, as new processing platforms emerge that overcome power and size constraints, it is very likely that SDR will make its way into portable devices. Indeed, some industry insiders are predicting that by 2015 there will be a transition from the current generation of handsets to SDR handsets.

1.3 COGNITIVE RADIO

1.3.1 What Is Cognitive Radio?

The reconfigurability offered by SDR technology enables radios to switch functions and operations. However, an SDR can do this only on demand; it is not capable of reconfiguring itself into the most effective form without its user even knowing it. In Mitola's dissertation and a number of publications, he envisioned such a

self-reconfiguring radio and dubbed the term *cognitive radio* for it. According to Mitola's early vision, a CR would be realized through the integration of model-based reasoning with software radio and would be trainable in a broad sense, instead of just programmable. In analogy with the mental process of cognition, Mitola also outlined a cognitive cycle through which such radio can reconfigure itself through an ongoing process of awareness (both of itself and the outside world), perception, reasoning, and decision making. The concept of CR emphasizes enhanced quality of information and experience for the user, with cognition and reconfiguration capabilities as a means to this end. Today, however, CR has become an all-encompassing term for a wide variety of technologies that enable radios to achieve various levels of self-configuration, and with an emphasis on different functionalities, ranging from ubiquitous wireless access, to automated radio resource optimization, to dynamic spectrum access for a future device-centric interference management, to the vision of an ideal CR. Haykin, for example, defines CR as a radio capable of being aware of its surroundings, learning, and adaptively changing its operating parameters in real time with the objective of providing reliable anytime, anywhere, and spectrally efficient communication [11]. The U.S. Federal Communications Commission (FCC) uses a narrower definition for this concept: "A Cognitive Radio (CR) is a radio that can change its transmitter parameters based on interaction with the environment in which it operates. The majority of cognitive radios will probably be SDR (Software Defined Radio) but neither having software nor being field programmable are requirements of a cognitive radio."

Despite these differences in both the scope and the application focus of the CR concept, two main characteristics appear to be in common in most definitions. They are *reconfigurability* and *intelligent adaptive behavior*. Here by *intelligent adaptive behavior* we mean the ability to adapt without being a priori programmed to do this; that is, via some form of learning. For example, a handset that learns a radio frequency map in its surrounding could create a location-indexed RSSI vector (latitude, longitude, time, RF, RSSI) and uses a machine-learning algorithm to switch its frequency band as the user moves.

From this it follows that cognitive radio functionality requires at least the following capabilities:

- **Flexibility and agility**, the ability to change the waveform and other radio operational parameters on the fly. In contrast, there is a very limited extent that the current multichannel multiradio (MC-MR) can do this. Full flexibility becomes possible when CRs are built on top of SDRs. Another important requirement to achieve flexibility, which is less discussed, is reconfigurable or wideband antenna technology.
- **Sensing**, the ability to observe and measure the state of the environment, including spectral occupancy. Sensing is necessary if the device is to change its operation based on its current knowledge of RF environment.
- **Learning and adaptability**, the ability to analyze sensory input, to recognize patterns, and modify internal operational behavior based on the analysis of a

new situation, not only based on precoded algorithms but also as a result of a learning mechanism. In contrast, the IEEE 802.11 MAC layer allows a device to adapt its transmission activity to channel availability that it senses. But this is achieved by using a predefined listen-before-talk and exponential backoff algorithm instead of a cognitive cycle.

1.3.2 Evolution of Cognitive Radio

The main precursors for CR research was the seminal work by Mitola and Maguire in 1999 and early spectrum measurement studies conducted as early as in 1995 to quantify the spectrum use, both in the licensed and unlicensed band. In the United States, CR research focused quickly on dynamic spectrum access (DSA) and secondary use of spectrum as the main objectives of the initial research. This was due to the fact that it was attracting a number of early research projects (e.g., URA, SPECTRUM, and MILTON). The most notable project in the spectrum management and policy research was the XG-project funded by DARPA. The main goal of the XG-project was to study the so-called policy servers and secondary-use technologies, particularly for military purposes. However, the early success of XG was pushing the community to study more broadly the possibilities of CR. Another boost for the research was given by several vociferous researchers (such as Lessig, Reed, and Peha), who pointed out that there are possible flaws in the current regulatory domain.

In the standardization domain, three major groups have emerged to work on relevant technologies and architectures: IEEE 802.22 and SCC41 (formally P1900) working groups and more recently ETSI's Reconfigurable Radio Systsems Technical Committee on CRs and SDRs. Also, the SDR Forum as an industry group has studied some CR-related issues. Commercially, the most advanced standardization activity is IEEE 802.22 and related research that aims to provide dynamic access to vacant TV spectrum. However, IEEE 802.22 requires a rather limited level of cognition.

At the time of this writing, CR is being intensively investigated and debated by regulatory bodies as the enabling technology for opportunistic access to the so-called TV white spaces (TVWS): large portions of the VHF/UHF TV bands that become available on a geographical basis after the digital switchover. In the United States, the FCC already proposed to allow opportunistic access to TV bands in 2004 [12]. Prototype CRs operating in this mode were put forward to the FCC by Adaptrum, I^2R, Microsoft, Motorola, and Philips in 2008 [13]. After extensive tests, the FCC adopted in November 2008 a Second Report and Order that establishes rules to allow the operation of cognitive devices in TVWS on a secondary basis [14]. Furthermore, in what is potentially a radical shift in policy, in its recently released Digital Dividend Review Statement [15], the U.K. regulator, Ofcom, is proposing to "allow licence exempt use of interleaved spectrum for cognitive devices" [15]. Furthermore, Ofcom states that, "We see significant scope for cognitive equipments using interleaved spectrum to emerge and to benefit from international economics of scale" [15]. More recently, on February 16, 2009, Ofcom published a new consultation providing further details of its proposed cognitive access to TVWS [16]. With both the United States and United Kingdom adapting the cognitive access model, and the emerging 802.22 standard for cognitive access to TV bands [17, 18] being at the final stage, we can expect that CR may become mainstream technology worldwide in the near future.

1.4 KEY APPLICATIONS

As discussed, CRs are highly agile wireless platforms capable of autonomously choosing operating parameters based on both prevailing radio and network conditions [11, 19]. Consequently, CRs have the potential to revolutionize how devices perform wireless networking. For instance, CRs allow radios operating on different protocols and standards to communicate with each other. This is known as *interoperability* [20, 21]. Furthermore, CRs are capable of transmitting in unoccupied wireless spectrum while minimizing interference with other signals in the spectral vicinity; that is, DSA [22–24].

1.4.1 Interoperability

Today, a plethora of wireless standards, applications, and services are being employed across numerous sectors of modern society, as well as within the same sector, such as military, public safety, and emergency responders [25]. Consequently, the use of multiple (potentially incompatible) communication standards within a specific sector could seriously impact the effectiveness of coordinated operations, yielding a situation analogous to the biblical account describing the Tower of Babel. For instance, the effectiveness of the emergency responders to cope with the aftermath of Hurricane Katrina in New Orleans during August 2005 was greatly affected by the inability of their diverse range of deployed communication equipment to operate with each other, especially within a decentralized operating environment.[1] This is shown in Figure 1.2, where members of Team A employ a communications standard operating on a carrier frequency that is different from the communication equipment employed by both Teams B and C. Thus, unless these teams are coordinated with respect to operating parameters and communication standards, effective communications between them would be nearly impossible.

Nevertheless, there are several reasons why sectors such as the military and public safety still maintain a range of communication solutions, such as significant financial investment and specific performance requirements. Consequently, CR possesses the potential to undo the Tower of Babel syndrome with respect to communications between teams employing different equipment [20, 21, 25, 26].

[1] Without reliable electrical power, cellular base stations and other centralized communication nodes were nonfunctional after their emergency power supplies were depleted.

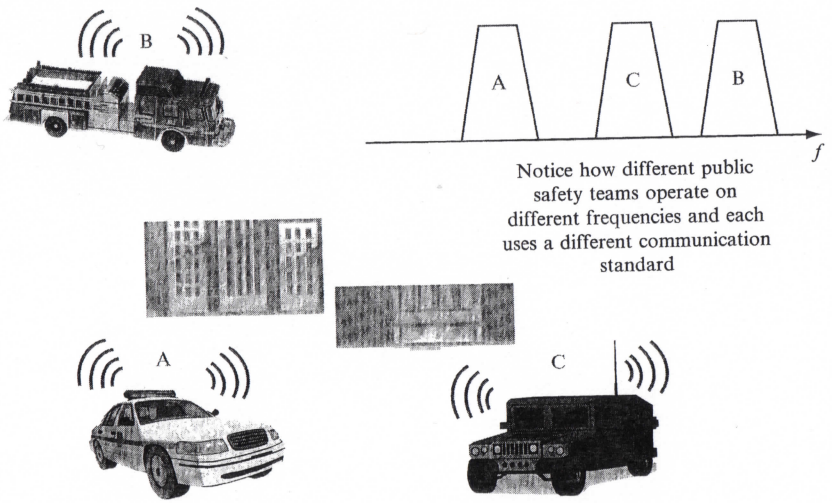

FIGURE 1.2

Example of public safety and emergency responder teams within the same geographical area operating on different center frequencies and potentially using different communication standards.

Due to its ability to rapidly assume any available radio configuration, CR platforms can reconfigure themselves to a legacy communications standard in order to communicate with any communication system deployed in the field or facilitate communications between two non-CR platforms employing different standards. Furthermore, with its onboard artificial intelligence, CR can automatically distinguish between different communication standards in the absence of any centralized control.

1.4.2 Dynamic Spectrum Access

With the increasing demand for additional bandwidth to support existing and new services, both spectrum policy makers and communication technologists are seeking solutions for this apparent spectrum scarcity. Meanwhile, measurement studies have shown that much of the licensed spectrum is relatively unused across time and frequency [1,27–32]. Nevertheless, current regulatory requirements prohibit unlicensed transmissions in these bands, constraining them instead to several heavily populated, interference-prone frequency bands. To provide the necessary bandwidth required by current and future wireless services and applications, the FCC has commenced work on the concept of unlicensed users "borrowing" spectrum from spectrum licensees [33,34]. This approach to spectral usage is known

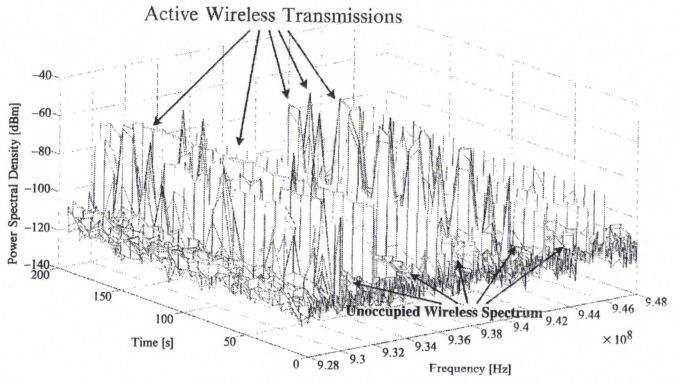

FIGURE 1.3

Wireless spectrum of 928–948 MHz in Rochester, New York, on June 19, 2008. Notice the time variations in the spectrum occupancy at several frequencies in this band.

as *dynamic spectrum access*. With recent developments in CR technology, it is now possible for these systems to simultaneously respect the rights of incumbent license holders while providing additional flexibility and access to spectrum.

An example of how DSA would work can be illustrated in Figure 1.3, where parts of the spectrum between 928 MHz and 948 MHz are occupied over both frequency and time. However, it is readily observable that there also exists portions of the spectrum that are unoccupied for a significant period of time, making them suitable candidates for secondary access by unlicensed wireless devices in a DSA framework. Nevertheless, when accessing these unoccupied frequency ranges within licensed spectrum, the secondary wireless device must ensure that it does not interfere with the operations of the primary user transmissions. Interference may occur when the out-of-band (OOB) radiation of the secondary transmission exceeds the tolerable levels, contaminating the primary user transmission if located relatively close in the frequency domain. Simultaneously, given the time-varying nature of wireless transmissions, a spectrum that might be unoccupied at one time instant could potentially be occupied at a subsequent time instant. Consequently, the CR platform must be environmentally aware and rapidly reconfigurable to prevent secondary user interference of primary user transmissions.

To achieve higher spectral efficiency, multiple access techniques can be employed such that multiple secondary users can transmit data within the same frequency range. Several techniques have been proposed to achieve multiple secondary user access, including those based on *code division multiple access* (CDMA) [35,36], spatial multiplexing [37], and *orthogonal frequency division multiplexing* (OFDM) [38,39]. With respect to OFDM-based techniques, the *spectrum pooling* concept can be effectively employed, where data are transmitted across unoccupied portions of frequency using a subset of active subcarriers [24].

1.5 BOOK ORGANIZATION

The chapters of this book have been grouped together into three thematically related parts to provide better structure for the reader with respect to the topics covered: Radio Communications, Networks, and Implementation; Applications; and Case Studies. These parts can be loosely divided into the following six subparts: The first subpart includes Chapters 1 and 2. It gives the reader some essential background of SDR and CR technologies and their impact on radio spectrum regulatory policies. The second subpart of this book includes Chapters 3 to 7 and focuses on the underlying physical layer technologies, spectrum sensing, reconfiguration, adaptation, and spectrum access, all of which are basic building blocks for a CR. In particular, Chapter 3 offers a review of the digital communications that are most relevant to the design of CRs. Readers who are familiar with this background may proceed directly to the other chapters.

The third subpart of this book consists of Chapters 8 to 12 and is centered around networking aspects of CRs. In particular, Chapter 8 offers a review of fundamentals of communication networks and may be skipped if readers already have this background. The fourth subpart of this book includes Chapters 13 to 17, which cover CR terminology, applications, and security issues. In particular, Chapter 14 offers an extensive discussion on how CR can share the spectrum with TV bands, while Chapter 16 discusses how CR can alleviate the interoperability problem for public safety communications. With respect to spectrum trading, Chapter 17 examines the challenges and solutions in this area with an emphasis on dynamic spectrum auctions.

The fifth subpart of the book includes Chapters 18 and 19 and focuses on CR testbed platforms. In particular, Chapter 18 gives an in-depth coverage of GNU radio and how to build a CR. Chapter 19 reviews some state-of-the-art testbed platforms for CRs. The last subpart of this book contains Chapter 20, which offers a perspective of CR architectural evolution and its future roadmap.