LEABHARLANN CHOLÁISTE NA TRÍONÓIDE, BAILE ÁTHA CLIATH
Ollscoil Átha Cliath

TRINITY COLLEGE LIBRARY DUBLIN
The University of Dublin

Terms and Conditions of Use of Digitised Theses from Trinity College Library Dublin

Copyright statement

All material supplied by Trinity College Library is protected by copyright (under the Copyright and Related Rights Act, 2000 as amended) and other relevant Intellectual Property Rights. By accessing and using a Digitised Thesis from Trinity College Library you acknowledge that all Intellectual Property Rights in any Works supplied are the sole and exclusive property of the copyright and/or other IPR holder. Specific copyright holders may not be explicitly identified. Use of materials from other sources within a thesis should not be construed as a claim over them.

A non-exclusive, non-transferable licence is hereby granted to those using or reproducing, in whole or in part, the material for valid purposes, providing the copyright owners are acknowledged using the normal conventions. Where specific permission to use material is required, this is identified and such permission must be sought from the copyright holder or agency cited.

Liability statement

By using a Digitised Thesis, I accept that Trinity College Dublin bears no legal responsibility for the accuracy, legality or comprehensiveness of materials contained within the thesis, and that Trinity College Dublin accepts no liability for indirect, consequential, or incidental, damages or losses arising from use of the thesis for whatever reason. Information located in a thesis may be subject to specific use constraints, details of which may not be explicitly described. It is the responsibility of potential and actual users to be aware of such constraints and to abide by them. By making use of material from a digitised thesis, you accept these copyright and disclaimer provisions. Where it is brought to the attention of Trinity College Library that there may be a breach of copyright or other restraint, it is the policy to withdraw or take down access to a thesis while the issue is being resolved.

Access Agreement

By using a Digitised Thesis from Trinity College Library you are bound by the following Terms & Conditions. Please read them carefully.

I have read and I understand the following statement: All material supplied via a Digitised Thesis from Trinity College Library is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of a thesis is not permitted, except that material may be duplicated by you for your research use or for educational purposes in electronic or print form providing the copyright owners are acknowledged using the normal conventions. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone. This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

Deconstructing and Reconstructing Medium Access Control: The Anatomy of a Cognitive Radio MAC

James Colman O'Sullivan

A thesis submitted to the University of Dublin
for the degree of Doctor of Philosophy

Department of Electronic and Electrical Engineering
University of Dublin, Trinity College

2012

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

Summary

This thesis focuses on Medium Access Control (MAC) for Cognitive Radio (CR). What is meant by the term can however be unclear. What CR MACs have been shown to do in the literature does not follow any one clear definition, nor is it based on any one understanding or set of principles. This variation and difference in understanding leads to difficulty in comparing solutions and identifying novelty. In short there appears to be no systematic way in which to *define* a CR MAC. Nor does there appear to be a systematic way in which to *design* a CR MAC. In addition there are limited examples of real working prototypes. While this in part may be due to a lack of interest on the part of individual researchers, it could also be concluded that the lack of systematic approaches mitigates against realisation of solutions.

The fundamental question asked by this thesis is how to create a means of systematically *defining* and *designing* CR MACs.

The core contribution of this thesis is to introduce the *Anatomy of a CR MAC*. The Anatomy provides a consistent and structured means of identifying the *elements* of a CR MAC in a manner that takes account of the special challenges of CR, such as additional dependencies on the flexible physical layer, and the additional functionality associated with the identification and sharing of spectral resources.

The Anatomy of a CR MAC facilitates design of CR MAC protocols with variable PHY behaviour and possibly multiple discrete PHY transceivers. It is a method that can be applied to any specific platform of implementation, and in this thesis, one such platform – Iris – is chosen as an example. The Anatomy varies from alternative mechanisms of dividing MAC functionality into state-machine like representations in that the elements of the Anatomy are of lesser granularity than any other modular representation, and represent unique and interesting functionality portions, rather than simple common parts which must be combined to provide meaning. The Anatomy requires that PHY-MAC overlap be noted in the form of PHY dependencies and contracts. It also distinguishes between the spectrum resource allocation and spectrum resource exploitation levels of a CR MAC. And in contrast with other

methods, it operates at the per-device level, taking a process-oriented rather than outcome-oriented view.

The Anatomy is put into action in several ways in this thesis. It is used to describe and chart existing CR MAC solutions. It is used to compare different CR MACs. It is used to combine elements from various MACs to create new CR MACs, and it is used to design a specific CR MAC from scratch, dubbed "CycloMAC". CycloMAC is subsequently implemented on a platform consisting of the Iris software radio and the USRP RF-frontend.

An enormous body of research has already been carried out in the area of CR MACs. The Anatomy of a MAC offers a mode of making sense of all of this work; a way through the literature; and as is mentioned and demonstrated a number of times in this thesis, a means of making use of the material in a structured manner.

The thesis begins by grounding the reader in CR and its concepts, the basics of MAC for wireless, and their union. It then presents an overview of selected works from the literature on CR MAC, summarises the issues and weaknesses with current CR MAC research, and introduces the challenges of defining and constructing an ideal CR MAC.

It introduces four key concerns that must be addressed in order to correct the issues and weaknesses of current CR MAC research, and to assist CR MAC research going forward. It then introduces the Anatomy as a novel means for addressing these concerns, and discusses its uses with examples. It builds on the introduction of the Anatomy by demonstrating how it can be used to inspire CR MAC implementations for real-world experimentation and use. Using the Iris Software Defined Radio (SDR) platform, it discusses the benefits of using the Anatomy as an implementation method. It details the design of a CR MAC – CycloMAC – tailored to a chosen scenario, and discusses its implementation.

The Anatomy is then analysed in greater depth, by discussing its use cases, both in general terms and with recourse to examples drawn from the literature and from earlier in the thesis. The thesis discusses the challenges the Anatomy faces, both in terms of representation of protocols and in assisting implementation. And it revisits related work, before concluding with a consideration of the future of CR MAC.

Acknowledgements

Heartfelt thanks are owed to a great many people who assisted me in the work leading up to this dissertation over the last four years, both directly and indirectly.

Firstly, deepest thanks are due to my supervisor Linda Doyle for her tireless enthusiasm, dedication to her students, and passionate interest in everything to do with wireless communication. Her unfailing ability to knit seemingly disparate pictures together and inspire others to regain perspective is really something special. A conversation with her almost inevitably reignited my excitement in my research, even at times when I'd convinced myself it couldn't possibly.

Thanks are also due to the many great researchers whose time at CTVR has overlapped with mine. To Tim, for providing wit, insight and an unrepayably large supply of coffee. To Paul, Barış and Keith for assisting me with Iris and spectrum analysers, and for reminding me that radio hardware does not necessarily bear a grudge. Thanks go to Jörg also, for help with LaTeX, for providing this template and for general technical-support.

And my thanks to everyone else in CTVR for the general good-cheer that makes it a pleasant place to work and to be.

Outside of the lab, I owe huge thanks to David for always being there, and for living in a world as far removed from telecommunications engineering as it's possible to be. To my mother Ann, and father Jim, thanks are due to both for their unwavering confidence in me. I'd like to thank Aisling for the supply of television that helped to make long days that little bit easier. Thanks to Declan, for never losing a sense of wonder about the world, and knowing exactly what it's like. To Ross, for unrelenting encouragement. To Alex, for counselling perseverance. And to everyone else, and in particular Pádraic and Alan, my thanks for providing oft-needed distraction, and for never asking me what my thesis was.

Finally, thanks are due to Tim, Dee and Séamas for their assistance in proofing this manuscript.

Contents

1 Introduction

The focus of this thesis is on Medium Access Control protocols for Cognitive Radio. In 1999 Mitola defined a Cognitive Radio as a context-aware radio device, that "knows what it knows","transform[ing] radio nodes from blind executors of predefined protocols to radio domain- aware intelligent agents that search out ways to deliver the services the user wants even if that user does not know how to obtain them" [1]. The definition of a Cognitive Radio (CR) has narrowed somewhat since that early definition. For the purposes of this thesis, a CR is a device bringing together a highly flexible and reconfigurable radio and some kind of decision-making mechanism(s) to create a communications system aware of its environment and capable of building up a knowledge-base, reacting to stimuli and learning from experience. While the definition of the CR can be clear, the definition of a Cognitive Radio Medium Access Control (MAC) fails to be succinct. The challenge of defining exactly what a CR MAC is, and the challenge of how Cognitive Radio MACs may be more systematically designed, are at the heart of this thesis.

This opening chapter briefly introduces the concept of the Cognitive Radio MAC, identifies key challenges in the field, enumerates the contributions of the thesis and maps out its remainder.

1.1 Introducing the CR MAC

Cognitive Radio has received significant attention in the past decade. In the main, a CR is seen as an entity that can facilitate dynamic access to spectrum. The concept of such Dynamic Spectrum Access (DSA) will be returned to in slightly more detail in Chapter 2. For the moment it is sufficient to mention that commercial, public safety and military scenarios that call for the dynamic accessing of spectrum have been envisaged and CR solutions for these scenarios have been explored. The working definition of CR used in this thesis is one that captures the flexible nature of an entity that must operate in conditions in which static allocations of spectrum do not exist. While the raison-d'être of a cognitive radio goes beyond DSA, the key point is that a CR reacts to the circumstances in which it finds itself, and does so sufficiently quickly to allow it consistently meet user performance requirements.

From a physical (PHY) layer perspective, a CR must be able to change its frequency of operation; manipulate its waveform characteristics including structure, bandwidth, modulation type; alter power levels and beam patterns; and many other properties, all to suit prevailing conditions.

In general terms, the purpose of a MAC protocol is to impose meaning on the raw bits transmitted by a communications device, and to ensure that where possible what is transmitted is also received. In the wireless domain, the radio frequency (RF) spectrum can be considered a shared resource to which devices have access. Simultaneous use of this resource can cause a failure in communications, therefore access to it must be mediated. Though there are only a small number of methods by which this resource can be divided, these methods can be combined and layered together in a variety of intricate ways, to create a significant number of wireless MAC protocols. On one level, this is exactly what a CR MAC protocol must also do. However, in traditional radio systems the spectrum resource to be shared and the sharing mechanism itself are well defined. This is not the case for Cognitive Radios. A CR MAC therefore should, in principle, involve itself in determining what spectral resources are available in the first place, decide how those resources might be shared out among users in the network or indeed between neighbouring networks, and then – as in the case of the traditional MAC – ensure that transmissions are received. It should, in principle, do all of this while the available resources and prevailing conditions change and it should, in principle, not cause interference to other networks and systems while carrying out its various functions.

1.2 Challenges

What CR MACs have been shown to do in the literature does not however follow any one clear definition, nor is it based on any one understanding or set of principles. This variation and difference in understanding leads to difficulty in comparing solutions and identifying novelty. In short, there appears to be no systematic way in which to define a CR MAC. Nor does there appear to be a systematic way in which to design a Cognitive Radio MAC. In addition there are limited examples of real working prototypes. While this in part may be due to a lack of interest on the part of individual researchers, it could also be concluded that the lack of systematic approaches mitigate against realisation of solutions.

The fundamental question asked by this thesis is how can we create a means of systematically defining and designing CR MACs.

1.3 Contributions

The core contribution of this thesis is to answer that question, through the introduction of a process that can be used to deconstruct CR MACs that have already been designed, and to construct new CR MACs in a methodical and structured manner. That process is known as the *Anatomy of a CR MAC*. The term anatomy is usually associated with the science of the shape and structure of organisms and their parts. Using similar vocabulary, the Anatomy of a MAC has been created to provide a systematic way to look at the shape and structure of the CR MAC and its different parts. The Anatomy provides a means of identifying the parts of a CR MAC in a manner that takes account of the special nature of the CR problem-space, such as the additional dependencies on a CR's flexible PHY layer, and the extra functionality associated with the defining and sharing of spectral resources.

The Anatomy addresses the challenges identified in the previous section. It aids in defining what a CR MAC really is. It provides a means of documenting the CR MAC solutions that already exist. It provides a means for recognising similarity and difference between varying MAC solutions. It provides a means of creating protocols by combining different parts of existing CR MAC solutions as well as providing a means of designing entirely new protocols. It informs real world implementation of CR MAC protocols. As part of the core contribution of the thesis, the Anatomy of a MAC is applied to these varying challenges. In particular, a fully implemented CR MAC designed using the Anatomy is presented.

Associated contributions include a detailed review of a selection of the wide body of literature in CR MAC and a recognition and explanation of the key aspects of CR MAC, gleaned from analysis of the literature.

1.4 Summary of this thesis

The thesis will begin in Chapter 2 by grounding the reader in CR and its concepts, the basics of MAC for wireless, and the intersection of the two. To contrast with the definitions of CR MAC it considers, it then presents an overview and analysis of selected works from the literature on CR MAC. The chapter will conclude by summarising the issues and weaknesses with current CR MAC research, and introducing the challenges of defining and constructing an ideal CR MAC.

Chapter 3 will expand on these issues, by introducing four key concerns that this thesis believes must be addressed in order to correct the issues and weaknesses as they relate to current CR MAC research, and to assist CR MAC research going forward. It will introduce the *Anatomy of a CR MAC*, a novel means for addressing these concerns, and discuss its uses with examples. It will conclude by making observations as to the effects of viewing CR MAC through the lens of the Anatomy.

Following that, Chapter 4 builds on the introduction of the Anatomy by demonstrating how it can be used to inspire CR MAC implementations for real-world experimentation and use. Using the Iris Software Defined Radio (SDR) platform, it discusses the benefits of using the Anatomy as an implementation method. Then it details the design of a CR MAC tailored to a chosen scenario, and discusses its implementation.

Chapter 5 analyses the Anatomy in greater depth, by discussing its use cases, both in general terms and with recourse to examples drawn from the literature and from Chapter 4. It discusses the challenges the Anatomy faces, both in terms of representation of protocols and in assisting implementation. And it revisits research from the CR MAC literature that involves similar thinking to that behind the Anatomy, as well as also discussing work not labelled as CR MAC research, but nonetheless of relevance. Finally, it uses these discussions as a springboard to reintroduce the topic of the ideal CR MAC, which it reconsiders in light of the Anatomy.

To conclude, Chapter 6 expands the scope of consideration to CR in general, in order to contemplate the greater context for CR MAC, before summarising and concluding the thesis.

2 A history of CR MAC issues

This chapter will explain the terms Cognitive Radio (CR), Medium Access Control (MAC), and MAC for CR. It will provide a definition for CR and discuss the ofttimes conflicting views of the technology given in the literature. As it unfolds, several important points detailing the contradictions and challenges in the field of CR MACs will be introduced. These include the lack of a commonly held and strict definition, and the general lack of and inconsistency of real-world implementation and experimentation. It will discuss the challenges of MAC for wireless devices in general, and detail how these challenges change and increase in magnitude in a CR. It will present an overview and analysis of the state-of-the-art in order to provide a firm grounding for the remainder of the thesis. Then finally, it will return to the points introduced throughout the chapter, and discuss how they motivate the remainder of this thesis.

2.1 Cognitive Radio

2.1.1 Defining a Cognitive Radio

Mitola defined a CR as a context aware radio device, that "knows what it knows", "transform[ing] radio nodes from blind executors of predefined protocols to radio-domain-aware intelligent agents that search out ways to deliver the services the user wants even if that user does not know how to obtain them" [1]. The decisions made by a CR might range from an attempt to make efficient use of radio resources by relocating to a frequency with less interference or accessible at lower cost, to reducing power consumption in order to conserve battery life, to being aware of a user's mental and physical distress in order to take actions of possible assistance [2]. Such a radio takes advantage of as many information sources as possible in order to inform its decision-making, from external sources such as satellite positioning and environmental sensors, to internal sources such as current operational efficiencies, Signal to Noise Ratio (SNR) and Bit Error Rate (BER) measurements, known frequency of operation and so on. In addition, it may sample and process the RF spectrum, not just to demodulate and to understand the signals around it, but to characterise and hence to further enhance its awareness of its spectral environment, and the behaviour of its neighbours.

2.1.2 Interpretations of Cognitive Radio

The broad definition of CR posited above is not necessarily one shared in the literature. The term is very often used in a much narrower sense, in which it is considered synonymous with *adaptive radio*, or occasionally Software Defined Radio (SDR)[1]. Adaptive radio is already a part of a great many successful and commercially available communications systems from cellular telephony to IEEE 802.11 "WiFi" [3] and IEEE 802.16 "WiMax" [4]. An adaptive radio is one that changes its operational parameters in response to some feedback. IEEE 802.11 & 16 for example change a signal's modulation order in response to measured SNR, in order to increase signal robustness in exchange for a reduction in data rate. Though this limited context-awareness might be thought of in the loosest sense as cognitive, it in reality is little more than a hardwired lookup table. A CR *learns* from its environment and experience, and may over time respond to the same situation in a different manner – hopefully with better results. Rita Mae Brown is thought to have written that insanity is doing the same thing repeatedly and expecting different results. If this is so, then merely adaptive radios are likely mad indeed [5].

This definition creep does lead to questions as to how "cognitive" a CR must be. What degree of cognition is "sufficient"? The label can be found in the literature applied to technologies as relatively simple as a fast frequency-hopping radio, or as advanced as a cognitive engine based on genetic algorithms and case-based reasoning such as that presented by Rondeau [6]. All this leads to a fractured body of research, which is based on inconsistent predicates and definitions. It should be pointed out that almost all research leans towards a simpler definition of CR, be it adaptive radio, SDR, or Dynamic Spectrum Access (DSA).

---

[1] A Software Defined Radio (SDR) is a radio in which the signal processing is performed where at all possible in software. This allows for a great deal of run-time flexibility relative to the fixed hardware pathways of conventional radios. SDR is a candidate technology for the implementation of a CR.

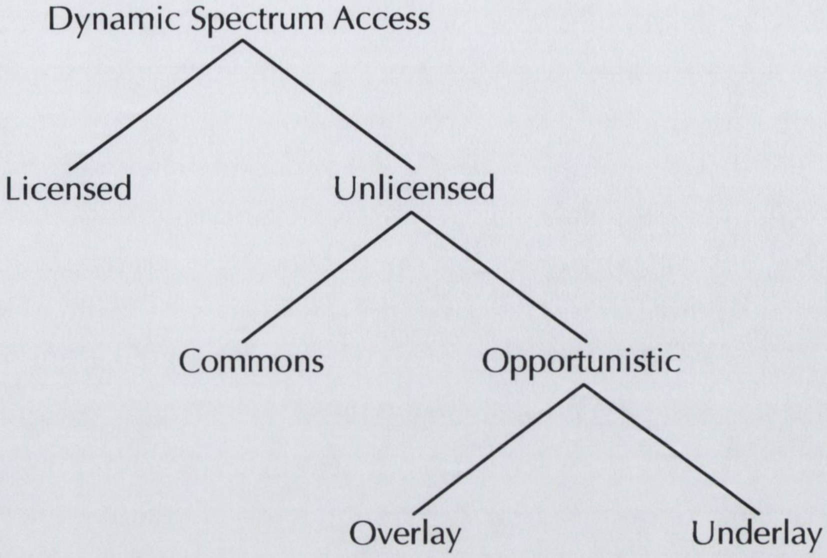

Figure 2.1: Tree representing the different modules of dynamic spectrum access.

2.1.2.1 Dynamic Spectrum Access

DSA is currently the most popular use case for CRs, and is a name given to the general idea of removing static spectrum assignment and moving to one or more of a number of more flexible access paradigms. These range from spectrum becoming "property" that can be bought, sold or rented; to unlicensed access where use of statically assigned spectrum by *secondary* users is permitted subject to legally defined and enforced constraints, and in particular the protection of the incumbent licensed *primary* users. It has come to the point that a reader of the literature may become confused as to where the distinction between DSA and CR is. This thesis will demonstrate that to many authors, there isn't one. As will be seen, commonly DSA is considered to be the application and a CR the radio employed as part of it, with the two inextricably linked. One could go further and point out that often even the breadth of interpretations of DSA is overlooked.

Given that, it's important to consider these interpretations. DSA may, as inferred above, be divided into one of two types, licensed and unlicensed. (See Fig. 2.1 for a depiction of the varieties of DSA.) Licensed DSA involves a change from the current model of static assignments, without the loss of the idea of of spectrum ownership. Unlike the current *command and control* model in which communications regulators license spectrum on a fixed and usually long-term basis, with no provision for transferability – indeed usually also mandating the specific technology to be deployed – licensed DSA would operate in an environment where permission to access the spectrum can be granted rapidly, perhaps in real-time. Licensees may be capable of reselling or leasing spectrum for short periods of time, in effect commoditising it. CR is an enabler for licensed DSA in that it allows radio hardware to change frequency and mode of operation to suit the spectrum its owner currently has access to. Ideally speaking, a CR would take responsibility for negotiating spectrum access directly, and be able to change spectrum rapidly and consistently. Its flexibility combined with knowledge of its own capabilities would allow it to attempt to act in the best interests of its owner, whether they choose to define that in terms of bandwidth, cost, power usage, reliability, or some combination [7].

Unlicensed DSA is itself divided into opportunistic and commons-based access. Opportunistic access offers a radical short term view of access and allows devices be permitted to access spectrum for which they are not the licensees, and without the licensees' permission, merely subject to certain legal constraints. This can function in one of two ways:

Firstly, the commonly cited *overlay* model permits devices make use of licensed spectrum when its licensee or *primary user* is not transmitting in it. This is often metaphorically described as "unoccupied" or "vacant" spectrum, and the individual moments can be referred to interchangeably as white spaces, spectrum holes, or spectrum opportunities. (See Fig 2.2.) Unfortunately such conceptually simple and pleasant adjectives bely the complex nature of the propagation of RF radiation, and might lead one to conclude that vacant spectrum can easily be defined and located. Techniques from the simple to the highly sophisticated have been proposed for defining and determining the presence of vacancies, including using radio hardware to "sense" them [8], and the use of databases of primary user activities. There remains considerable disagreement both legal and technical as to the level and method of protection required for primary users.

Secondly, an alternative model for opportunistic DSA is that of *underlay*. This involves secondary users taking advantage of spread-spectrum communications technologies, in order to spread their signal over a large bandwidth but at a reduced power, thus allowing safe transmission simultaneously with primary users, up to a regulator specified interference threshold [9].

In contrast to opportunistic access, a commons-based DSA model is one in which no device has a license, but instead all devices have a "right to transmit" [7]. This contrasts with conventional licensing, in which licensees have an implicit "right to receive." In a spectrum commons, CR devices may employ an etiquette or rule based access to ensure reasonable mutual fairness and use of the spectrum [10].

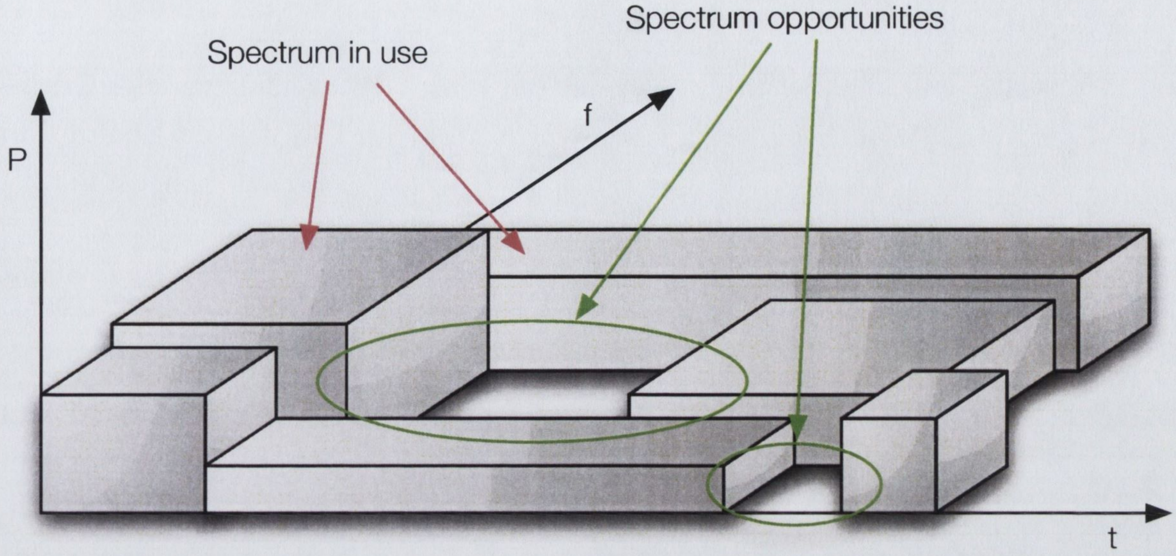

Figure 2.2: A simple representation of overlay DSA. Spectrum not in use by a primary user can be *re-used* by a secondary user, subject to agreed constraints.

A diagrammatic representation of the relationships between the different models of DSA was shown earlier in Fig. 2.1. However, though these DSA models may be technological contributions, they remain just one application of a CR, a fact which must be remembered when reading literature which may consider CR and DSA to be synonymous.

2.1.3 How this thesis defines Cognitive Radio

Contrasting two of the definitions for CR discussed thus far, the first is Mitola's, of an extremely flexible context aware radio device which optimises user experience and access efficiency, taking into account all possible context, from the technological to the social. The second, narrower definition is unlicensed DSA, a radio that makes use of spectrum allocated to other devices, in a manner that minimises impact to them. That is not to say that unlicensed DSA is not an interesting scenario, and potentially one in which a highly flexible radio can be brought to bear, but it nonetheless narrows scope relative to the broadest possible definition of a CR.

This thesis accepts a definition that lies in the middle ground – that it is a device bringing together a highly flexible and reconfigurable radio and some kind of decision-making mechanism(s) to create a communications system aware of its environment and capable of building up a knowledge-base, reacting to stimuli and learning from experience.

This thesis considers Medium Access Control in CRs. Therefore with CR introduced and defined, in order to best explain MAC for CR, it is next helpful to introduce MAC for wireless devices in general. The next section will do just that.

2.2 Medium Access Control for wireless

A MAC protocol is responsible for assigning order to the raw bits received and transmitted by the physical (PHY) layer. These bits correspond to little more than changes in voltage in the medium – so it is the MAC protocol which gives them meaning. The transmission becomes a *message*. At its basic level, a MAC protocol is really quite simple. Considering wireless communications specifically, each device has access to an radio frequency (RF) *spectrum resource* determined by its PHY. The MAC mandates when and where transmissions should take place within this spectrum resource. There are only a small number of ways this resource can be shared – in time, in frequency, in space, and in a logical sense in code.

A MAC protocol's responsibilities can be expressed in a few short points, in order of importance:

1. Choosing where and when to access the spectrum resource.
2. Ensuring message delivery.
3. Arranging necessary control signalling.
4. Handling the edge-cases[2] of the chosen access mechanism.

These responsibilities are arranged in order of their importance, and in fact simple MACs may only meet the first of them.

Note that from now on and for ease of discussion, this thesis will refer to "a MAC", "MAC layer" and "MAC protocol" interchangeably.

2.2.0.1 A note on the scope of this definition

It should be noted at this point that MACs in most modern communications systems are thought of in a highly static light. They are designed to fit together with a given PHY, and usually reside on the same silicon. Therefore it is common that a given technology such as Bluetooth [11] or IEEE802.11a [3] will be thought of and

---

[2] This term will be defined in §2.2.1.4

referred to as a wholly contained MAC/PHY pair. Essentially this link is borne of implementation choices and the heretofore static behaviour patterns of communications protocols. In the definition presented above and expanded below, this thesis seeks to break this automatic MAC/PHY connection, as in discussing Cognitive Radio it considers MAC for a technology that is not fixed in the way of conventional standards-based communications devices. As a result, the responsibilities of a MAC presented above must be considered as abstract from a specific PHY layer.

2.2.1 Basic Medium Access Control issues

Though the basic issues of Medium Access Control for wireless are well understood, this section will discuss them in greater detail for the completeness of this thesis, and in order to lay a consistent theoretical and terminological framework for the reader.

2.2.1.1 Choosing where and when to access the spectrum resource

At its simplest level, the problem any MAC protocol attempts to solve, whether wired or wireless, is to prevent or minimise simultaneous use of the medium occurring. In both wired and wireless systems, an enormous increase in the understanding of signal processing in recent decades has refined what is meant in this context by the word *simultaneously*.

Telecommunications engineers now use it to mean transmissions centred or overlapping on the same frequency, reasonably close in space[3] and at sufficient power level to disrupt successful reception to an extent defined unacceptable.

In considering MAC design, an initial issue is whether the protocol will be centralised or distributed. A centralised MAC protocol usually offers superior performance and more predictable behaviour given that decisions are made with full knowledge of the topology of the network. On the other hand, a distributed MAC protocol is more flexible and resilient, and less likely to require some form of infrastructure.

Should a MAC fail to prevent simultaneous transmission, signals that are sent on overlapping frequencies, from devices that are sufficiently close, will *collide*. They will interfere, and any number of factors including – but not limited to – geography

---

[3] In the case of wireless communications, "close" here is a function of the propagation distance of the signal and the receiver sensitivity of the radio. Meanwhile, propagation distance is itself a function of transmission power, frequency and geography.

and hence path of transmission, power, modulation scheme and time of sending, may lead to the signal being unsuccessfully detected and demodulated at the receiver. This is obviously undesirable as it represents both a waste of energy and time. It should be noted that the interference does not necessarily manifest as a precipitous drop in communications effectiveness, but rather a gradual increase in the frequency of disruption. Furthermore, PHY techniques such as Forward Error Correction (FEC), or a switch to a more resilient modulation order, can be employed to mitigate this disruption to an extent.

Even with such PHY techniques employed to mitigate interference, there comes a point at which it becomes intolerable. Therefore, to prevent this, the RF medium must be divided in some manner and allocated to devices for safe transmission. Frequency can be channelised, time can be slotted, power can be controlled and space can be divided into cells. All of these divisions may be fixed or changing, and discrete or continuous in terms of their relative sizes.

In summary, given an RF spectrum resource, the wireless MAC must decide how to make use of it, and when each device should do so, in order to minimise interference.

2.2.1.2 Ensuring message delivery

A MAC protocol is also usually responsible for guaranteed delivery of data. Therefore it must implement a scheme to retransmit or recover data should undesirable collisions occur.

The most common scheme is Automatic Repeat reQuest (ARQ), in which a sending device will wait for a pre-agreed period of time after transmission for a small acknowledgement (ACK) message from its destination to indicate successful reception. If this ACK is not received, the transmitter will re-send its data, thus messages are assumed to have failed until notice is given to the contrary.

In order to make use of ARQ, it is also necessary to be able to determine whether the message received is identical[4] to that which was transmitted. Cyclic Redundancy Check (CRC), parity and/or other error detection mechanisms are used to allow a MAC protocol to detect errors in reception with a high degree of accuracy [12].

Integral to ensuring message delivery is the ability to specify a destination. RF is effectively a broadcast medium. Therefore when sharing a channel, every device

---

[4] Or as close as can reasonably be determined to identical.

within range of a transmission can receive it, whether it was the intended "destination" or not. A MAC protocol must incorporate some form of addressing mechanism so that devices can ignore transmissions that are not intended for them. This mechanism is also likely to require the capability to indicate that a transmission is in fact intended to be received by all devices, termed a broadcast address.

2.2.1.3 Arranging necessary control signalling

In order to meet its responsibilities, a MAC protocol is extremely likely to need to employ a certain amount of control signalling. Control messages may be needed to arrange transmissions, confirm successful reception, notify devices of resource assignments, change resource assignments and perform any one of many small tasks. A MAC protocol must include the structure of all control communications as part of its specifications, in addition to the framing placed on data messages.

2.2.1.4 Handling the edge-cases of the chosen access mechanism

The choice of how to divide and allocate the medium made by a given MAC protocol can often have edge-cases. An edge-case occurs when one or more of the assumptions which a MAC protocol makes is violated, and causes either disruption to communication, or unnecessarily conservative resource allocation. Carrier Sense Multiple Access (CSMA)-type MAC protocols suffer from two of the most commonly known edge-case issues, the *hidden node* and *exposed terminal* problems, which will be explained in the following section.

As noted earlier, these responsibilities of a MAC are arranged in order of their relative importance. It should therefore be pointed out that a great many protocols do not address the edge-cases of their access mechanisms. Doing so is not a strict requirement, but rather a desirable action, which will lead to increased performance and reliability.

2.2.2 Common choices for wireless MAC

In the broadest sense MAC protocols may be divided into two categories. These are scheduled protocols and contention-based protocols.

Scheduled protocols assign a transmission schedule to devices in advance of transmission, usually valid for a period of time. Depending on the nature of the mechanism determining the scheduling, these may sometimes be referred to as deterministic protocols.

Contention-based protocols – which are a variety of non-deterministic protocol – are the alternative, introduce a degree of randomness and demand-driven performance to their behaviour.

Both classes of protocol have their benefits – scheduling protocols are generally simpler to understand and analyse, and offer more predictable performance. Contention-based protocols on the other hand usually handle significant variations in levels of demand for access to the medium much better. Depending on the communications scenario, one, other or a hybrid approach combining both may provide the best performance, however that has been defined.

Probably the simplest and one of the earliest wireless MAC protocol is Aloha. Aloha is a non-deterministic protocol in which a device that has data to send simply sends it immediately. It incorporates an ARQ, and if required to re-send, the transmitter will wait a random period of time before doing so, called a *backoff*. This backoff helps to avoid subsequent collisions between transmissions. The choice of algorithm for calculating the backoff period can have a significant effect on the overall performance of an Aloha system [13].

A significant advance is to "listen-before-talking", receiving on the medium for an agreed period of time, and only transmitting if no signal from another device is detected. In order to attempt to be certain nothing is transmitting, this period of time is usually the maximum propagation time a message could take to pass between two devices. Provided the duration of data packets on the air is significantly larger than this maximum propagation time, listening before talking offers a marked throughput performance improvement over Aloha. The combination of listening before talking with a backoff in the event of a collision is called Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). This protocol has the advantage of preventing the majority of the collisions that would occur in the case of an Aloha protocol, and is almost undoubtedly the most popular of wireless MAC protocols. It should be noted that its applicability is limited by the requirement that packet sizes be appreciably larger than the propagation time. Should devices be spaced sufficiently far apart, the time required to sense may exceed the expected time to transmit the message, becoming wasteful of the medium. This limitation restricts CSMA/CA techniques to relatively low powers and/ or distances [14].

Note that when the medium is a wire, being aware of the occurrence of collisions is somewhat easier than when it is RF. Over a wire, should two or more devices transmit simultaneously, the event can be detected by those devices by simply noting a change in the received power on the line which does not correspond to their own transmissions alone. This is *collision detection* and is employed in a great many wired MAC protocols. In a wireless medium these collisions – where two or more transmissions destructively interfere – happen at the receiving device and so cannot be directly detected by a transmitter. Therefore a wireless MAC protocol must attempt to circumvent this limitation [15].

A hidden node occurs when a device attempts to communicate with a second, and successfully senses the carrier as idle. It starts to transmit to its intended destination, whilst unbeknownst to it, the second/ destination device is already receiving from another device which is sufficiently distant that the transmitting node cannot detect it. This third device is "hidden" to attempts by the first device to sense it. Should the first device transmit, the destination will be unable to receive either transmission intended for it. In effect, this is a collision which cannot be prevented, owing to the fact that interference is manifested at the receiver rather than at the transmitter [16].

The exposed terminal problem occurs by contrast when a device wishing to transmit senses the carrier busy, but the transmission it senses is to a device out of its own transmission range. Therefore if it were to transmit, there would be no collision at the currently transmitting device's receiver, and both transmissions would be successful. Instead the results of carrier sensing will cause the device to conservatively hold off on transmission until its neighbour is finished transmitting. This wastes time and reduces overall throughput [17]. See Fig. 2.3 for an example of both.

To borrow a succinct description from Brodsky *et al.* [18]: when a node is hidden, devices that should not transmit simultaneously do, and when exposed, devices that could transmit simultaneously don't.

CSMA/CA can further be enhanced by use of Ready To Send (RTS) and Clear To Send (CTS) messages, which help to alleviate the hidden node problem. Before sending a DATA message, and after sensing the carrier idle, a transmitter sends a short RTS message to its intended destination. Assuming the carrier is idle at the receiver also, the RTS will be successfully received, and the receiver will respond with a short CTS message. Then DATA and Acknowledgment (ACK) message transmission occur as usual. This RTS-CTS handshake brings multiple benefits. Firstly, it mitigates the hidden node problem: Neighbouring devices that receive both these

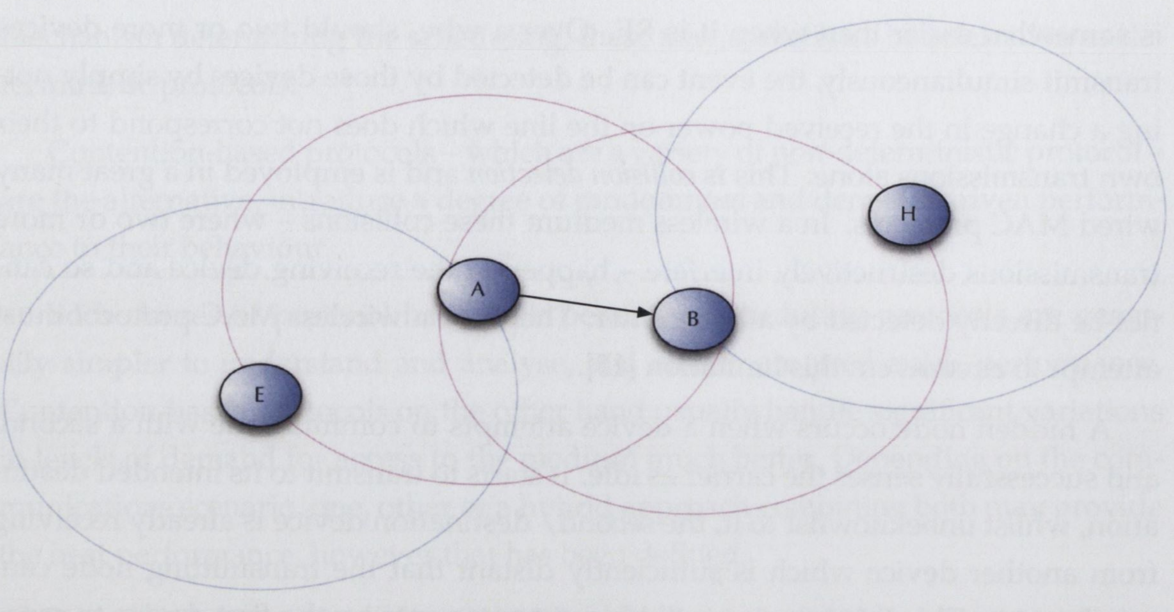

Figure 2.3: A graphical example of the hidden node and exposed terminal problems. Device *A* is transmitting to device *B*. If node *E* wishes to transmit to a device which is out of the range shown in this diagram – to the left, perhaps – it will unnecessarily defer its transmission, as *A*'s transmission leads it to sense the carrier busy. It is an exposed terminal. Meanwhile, node *H* might proceed with its own transmission to *B*, or to another node. It will incorrectly sense the carrier to be idle, and interfere with node *B*'s reception of signals from node *A*. It is a hidden node.

messages as they're exchanged can deduce they are within interference range of both devices. But a device that hears only a CTS message from a neighbour can infer that it must be hidden to the device transmitting to that neighbour. It will defer transmission until it receives the neighbour's ACK, or an appropriate timeout is reached. Additionally, a device which receives an RTS but not a CTS can infer that it is either an exposed terminal, or that the CTS was not sent by the neighbour's intended destination, and in either case can proceed to consider the medium to be clear. Though they bring additional overhead in successful packet exchanges, RTS-CTS handshakes can help reduce overhead when collisions do occur. In these cases, collisions will usually occur between two RTS packets. As these may be an order of magnitude smaller in size than DATA packets, the time taken for the intending transmitters to note the collision (by means of a missing CTS packet) and hence to initiate backoff is significantly reduced. Finally, this handshake also allows for the introduction of virtual-Carrier Sensing, which increases energy efficiency. If a transmitting device includes information as to how long the medium will be occupied by the intending transmission – called a Network Allocation Vector (NAV) – it can simply refrain from sensing the medium for that period of time. Therefore RTS and CTS messages include this information. The most common example of this variety of protocol is the Distributed Coordination Function (DCF) of the IEEE 802.11 wireless LAN specification [3].

The alternative to these contention-based and non-deterministic MAC protocols is *scheduled* MAC protocols, the most common of which is the Time Division Multiple Access (TDMA) family of protocols. TDMA protocols divide time into slots, which are agreed units of time, and assign slots to devices, usually for their exclusive use. This slot assignment for a group of devices is called a schedule. A group of slots is usually referred to as a frame, and will repeat at a guaranteed time interval. Scheduled protocols allow throughput guarantees to be offered to devices, and hence a potentially greater degree of reliability for devices, a quality that is highly desirable for many forms of communication, in particular 2-way voice and/ or video.

TDMA protocols differ in terms of their slot lengths, how assignment occurs, whether assignment changes over time, whether is made by a central entity or by some form of distributed consensus, how and what control signalling is exchanged, and so on.

Another option for MAC is Code Division Multiple Access (CDMA). Devices employing CDMA protocols use spread spectrum PHY layers, in which their intended signal is spread over a much greater bandwidth, than that strictly necessary. The three most common techniques for this are Direct Sequence Spread Spectrum (DSSS), Frequency Hopping Spread Spectrum (FHSS) and Time-Hopping Spread Spectrum (THSS). All employ a pseudo-random sequence or *spreading code* in order to spread the signal. A DSSS device spreads its signal directly by XORing it with the spreading code, while maintaining constant overall power. A FHSS device uses the spreading code to control short frequency hops among pre-defined channels. A THSS device employs the code to control direct access to the medium, with the code switching on and off the signal [19]. In each case, multiple CDMA devices transmit simultaneously, using spreading codes carefully chosen to be orthogonal. Provided the spreading codes are known at the receiver, the intended signal can be extracted from all the received signals. CDMA protocols require particularly tight integration between the MAC and PHY layers of a device, as codes must be known and shared, and precise power-control is required in order to minimise inteference within the system, called Multiple Access Interference (MAI). DSSS CDMA is also particularly vulnerable to the *near-far problem* in which a receiver is desensitized to weak signals

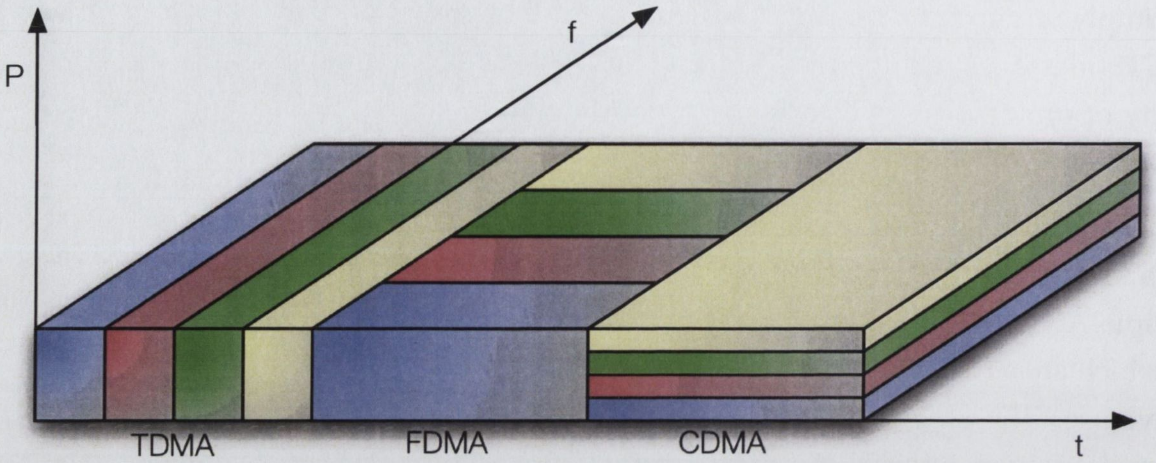

Figure 2.4: A simple graphical representation of TDMA, FDMA and CDMA, for the same block of frequency at different points in time.

by the presence of a strong signal. The assignment of codes allows CDMA protocols to in a logical sense be considered scheduled protocols [20].

The third model for scheduled MAC protocols is Frequency Division Multiple Access (FDMA). In an FDMA MAC protocol, devices or links are assigned unique frequencies for their use, usually referred to as channels. The model for assignment, and how it is exchanged with devices is the challenge here, as once assigned, frequencies are wholly reserved for the use of one device or one link, and so there is no risk of collision. This is of course a curiously idealised situation and is unrealistic both as it is rare for two devices to simply communicate with one another and no-one else, and also as sufficient bandwidth is unlikely to be available to a device that whole blocks of frequency can be assigned to devices on a permanent basis. As a result, FDMA is often incorporated as part of MAC protocols, for example in Global System for Mobile communications (GSM), which employs FDMA and TDMA.

All three scheduled protocol types are represented in simple conceptual form in Fig 2.4.

Most scheduled protocols take into account the effect of signal degradation over distance, and so take advantage of frequency, slot and code reuse for devices that are sufficiently far apart. This is sometimes referred to as Space Division Multiple Access (SDMA), but this is something of a misnomer as it – not unlike FDMA – is incorporated into other protocols, but does not exist on its own.

Of course given there are benefits to both scheduled and contention based protocols, it is no surprise that a combination of features from both is a common choice.

Slotted Aloha is one such, combining timeslots like TDMA and the random access benefits of Aloha. In this protocol, devices can only commence transmission at the beginning of a centrally controlled timeslot. This reduces the likelihood of collision by coordinating start times, and hence offers increased overall throughput.

More advanced MAC protocols differ primarily in their choice of how to perform resource assignments, and the control messages and the degree to which they perform precise per-transmission assignments that mitigate edge-cases as well as possible.

The discussions in this section for the most part describe what is often termed a single-channel MAC. Each device communicates on the same frequency, and all devices in a neighbourhood can receive one another's transmissions. FDMA represents an exception to this, as it allows devices to communicate concurrently on separate frequencies. The example of GSM was provided, in which FDMA and TDMA are employed. It would be more accurate to say that FDMA is employed in a static or very slowly-changing sense, while TDMA represents the message-to-message MAC protocol of GSM networks. This is because any given GSM base-station – and hence its surrounding cell – is assigned one channel, which it divides among its connected devices using TDMA. A connected device will only change channel when moving from the coverage of one base-station to another.

Frequency assignments in GSM are made with recourse to the assignments of other nearby base-stations, and to transmission power. No base-station's cell at frequency f_n will border on a the cell of a base-station also using f_n. Considering a base-station and all its connected handsets as a group of devices all within transmission and reception range of one another, a neighbourhood, there is no strict benefit to it in changing the frequency it uses, as all other GSM frequencies are occupied by neighbouring cells[5]. If these frequencies were not statically assigned however, and frequency assignments could be changed easily and quickly, new communications options might arise. Multiple communications between devices in this neighbourhood could occur simultaneously. Such a MAC is termed a *multi-channel* one.

2.2.3 Multi-channel MAC

Multi-channel MACs combine FDMA on a message-to-message timescale with some other access control mechanism. Using a multi-channel MAC protocol, multiple

---

[5] These neighbouring cells may not be within the range necessary for successful transmission and reception, but are likely to be within interference range, which is usually significantly larger.

devices all employing the same protocol can be communicating concurrently in separate channels, making use of them according to demand. Initiation and control of communication among devices, and ensuring no simultaneous communication occurs in any of the channels becomes significantly more difficult than in the single-channel case. It introduces three new challenges, *channel allocation*, the *multi-channel rendezvous* problem, and the *multi-channel hidden-node* problem. Furthermore, the broadcast of messages to all nodes which occurs by default in a single-channel MAC, must be explicitly arranged in a multi-channel MAC.

2.2.3.1 Channel allocation

Channel allocation is the actual choice of channel for each communication. It is an important differentiating factor between multi-channel MACs, and the choice may be made on the basis of anything from neighbour assignments and complex frequency re-use and power control models, to a simple random decision.

2.2.3.2 Multi-channel rendezvous

The multi-channel rendezvous problem is the name given to the difficulty in arranging communication with another device, when its current frequency of operation may not be known. Arranging communication with another device generally presupposes that the sender is capable of sending a control message to it to engage in a handshaking procedure – if there are a large number of possible channels available, there is no guarantee as to which one the receiver will currently be tuned to. There are three common classes of solution to this problem [21].

1. *Common Control Channel (CCC)*: This is a dedicated channel for control messages, on which data communication on other channels is arranged. In order to be as effective as possible, protocols making use of a CCC usually require devices to be in possession of a transceiver dedicated for that purpose. (See Fig. 2.5a.) In the event of the control and data communications taking place on a shared transceiver, the device is vulnerable to the multi-channel hidden-node problem.
2. *Split Phase*: Time is slotted, and every device communicates on a pre-agreed control channel for a control phase at the start of a slot in order to arrange communication on the other channels for the remainder of the slot. This solution requires devices to be time synchronised, which adds to device and/ or

protocol complexity. In certain cases, this can be referred to as an in-band CCC. (See Fig. 2.5b.)

3. *Hopping Based*: There are multiple sub-variants of this, but each involves devices frequency hopping between available channels, and using shared knowledge of the hopping sequences in order to arrange and perform communication. Depending on the choice of hopping solution, multiple devices may be capable of rendezvousing simultaneously. Hopping also requires devices to be time synchronised, in order to align hops (See Figs. 2.5c and 2.5d, which depict the common-hopping, and default-hopping mechanisms, respectively).

Note that these three classes are not mutually-exclusive. A split-phase control mechanism usually operates on a CCC for example, and a hopping-based mechanism could employ a common control-phase followed by an independent hopping period.

2.2.3.3 Multi-channel hidden nodes

The multi-channel hidden node problem occurs when a device transmits in ignorance on a channel which another has currently reserved for communication. It is caused by incomplete information on local neighbourhood activity, and is itself an edge-case caused by some solutions to the multi-channel rendezvous problem. Only a multi-channel rendezvous solution in which all nodes directly listen to the handshaking – and hence channel assignment – of neighbourhood devices at all times avoids this problem. A CCC for which each device has a dedicated transceiver, or a split phase control scheme both represent reasonable options for mitigation of this problem. In other cases, additional neighbour notification schemes may be necessary to mitigate or solve it. Note that even when the multi-channel hidden node problem is solved, the single channel hidden-node and exposed-terminal problems can still occur within each individual channel, if carrier sensing is being used.

Other more complete solutions include using an extended handshaking mechanism in which neighbours send warning messages to advise of channel occupancy in response to the RTS-CTS exchange of their neighbours; extending the range of the control channel to twice that of the standard data communications channels in order to guarantee RTS-CTS exchanges are heard; or use of Busy-tone(s) such as in [22].

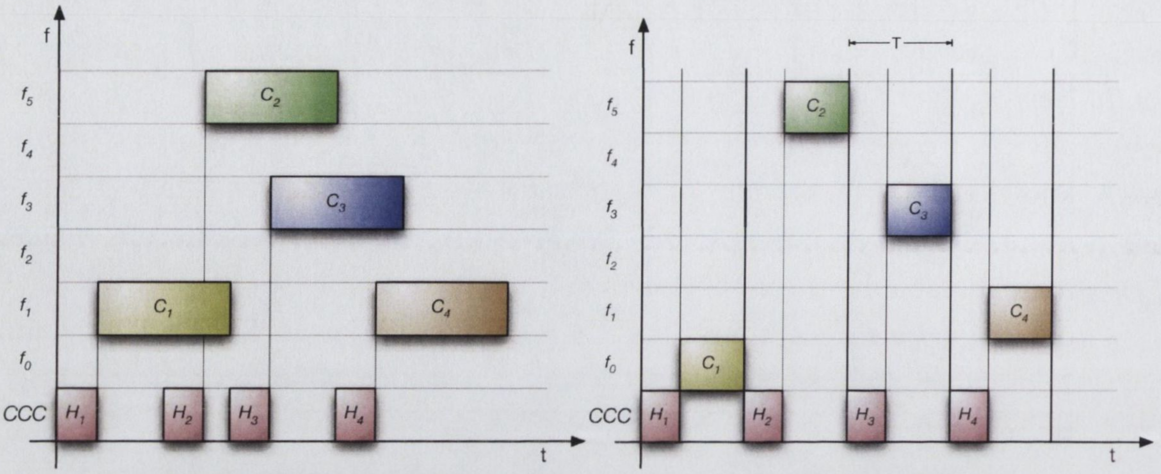

(a) Common Control Channel

(b) Split Phase

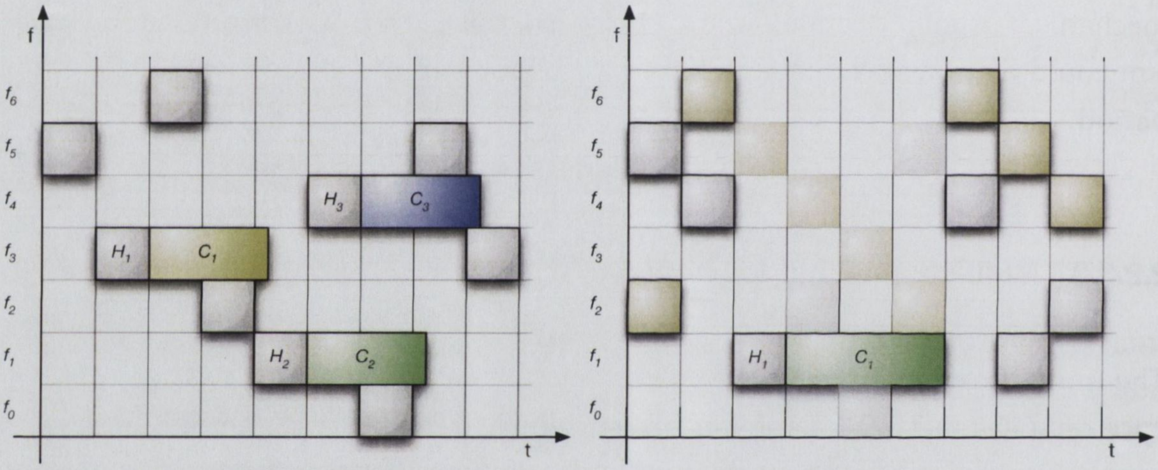

(c) Common Hopping Sequence

(d) Default Hopping Sequence

Figure 2.5: The four common multi-channel MAC control mechanisms. In each case H_i represents a packet-exchange handshake, such as an RTS-CTS exchange, between devices establishing communication. C_i represents the corresponding sequence of communications.

2.2.4 Duplexing

Not strictly discussed up until now is the subject of duplexing, or the control of bi-directional communication. As simultaneous transmissions from multiple nearby devices cannot usually be understood by a receiver, it follows that a device cannot transmit at the same time it receives, at least on the same frequency. If its own transmissions are held until it has finished receiving, a device is said to employ Time Division Duplexing (TDD). Alternatively, if it has the necessary PHY, it may transmit on one frequency and receive on another simultaneously, which is called

Frequency Division Duplexing (FDD). CSMA/CA MACs are usually TDD, while the example of GSM, and many other infrastructure based networks use FDD, and refer to the separate frequencies as the uplink (subscriber-device to base-station) and downlink (base-station to subscriber-device) respectively.

2.2.5 Efficiency & Performance

In spite of the increases in efficiency over the years, most wireless MAC protocols are enormously inefficient. Of the raw physical throughput offered by a given channel, bandwidth, modulation scheme and power level, even a "good" MAC protocol may only come close to a small utilisation of the possible throughput. The IEEE 802.11n MAC for example, when employing a 20MHz channel and a packet length of 1200B offers an effective throughput of only 21% of the raw PHY throughput [23]. TDMA protocols offer a much higher utilisation, but this assumes that all devices make full use of their respective timeslots, and as the transmissions of most consumer devices are bursty in nature, this may lead to wasted timeslots – one device may have a large amount of data to send, but is prevented from doing so in all timeslots not its own, whereas other devices may not be using their own timeslots. Many schemes for re-assignment of slots and the temporary use of empty slots have been proposed, and hybrid schemes combining scheduled and contention-based mechanisms such as Packet Reservation Multiple Access (PRMA) [24] have been devised. Nonetheless, such schemes all involve a significant level of overhead relative to the raw bandwidth available from a given spectrum resource, and furthermore tend to perform best in specific scenarios.

2.2.6 Recent innovations in wireless MAC

Notwithstanding the outline of MAC for RF wireless communications just described, some recent innovations have demonstrated that there remain boundaries that can be stretched, and fundamental assumptions that while not breakable outright, can be sidestepped to a surprising degree.

Choi *et al.* and Duarte *et al.* both recently proposed and described working experimental models of wireless links which allow full-duplex communication on the same frequency [25, 26]. By employing a variety of antenna based cancellation techniques with analog and digital domain signal cancellation mechanisms, both managed to demonstrate systems that can attenuate the self-interference caused by transmitting while receiving by between 60 and 80dB. This effectively renders it a non-issue, allowing a device to transmit while receiving. The resulting systems would be capable of such unconventional actions as: acknowledging safe receipt of messages *while* they continue to be received, forwarding a message in a chain of messages while receiving subsequent ones, and sensing the carrier even while idle, potentially reducing the risk of a collision at a third party device.

Another recent and not unrelated innovation is proposed by Sen *et al.*, who also provide a means to sidesteps an oft-cited law of wireless MAC protocols, that collisions cannot be detected by a transmitter [27]. By employing sequences known at the transmitter and receiver, a receiving device employing their mechanism will monitor hints as to received signal integrity from the PHY transceiver, and should a signal be deemed likely to have failed, it will transmit the known sequence, without waiting for the collided transmission to finish. A second antenna at the receiving device listens for this sequence, cross-correlating what it receives with the known sequence, in order to detect its presence in spite of self-interference from its own much stronger transmission. In order to increase the effectiveness of this correlation, digital domain signal cancellation techniques similar to those employed to enable full-duplex communications by Choi & Duarte are employed [25, 26]. This allows a transmitter to be aware of the occurrence of a collision prior to the completion of the unsuccessful transmission. It can therefore abort and retransmit immediately. Such a mechanism largely removes the need for backoff mechanisms, offering enormous potential performance improvements for CSMA type protocols.

These systems represent a possible first step in a changed model of medium access control, but they do not remove all limitations on access to the medium, and leave remaining a host of potential possible collisions. Furthermore, they suffer from a number of hardware limitations, including significant processing requirements, the need for multiple physically separated antennas and strong frequency specificity.

2.2.7 Blurred boundaries

Of course this discussion has treated MAC protocols for the most part as separate from the PHY layer. Unfortunately, as mentioned earlier in §2.2.0.1, in real radio devices this boundary is not so clear. In order to make effective decisions and offer timing guarantees, MAC protocols for currently available devices are almost invariably implemented in hardware, and usually fixed-path hardware such as an

Application Specific Integrated Circuit (ASIC) at that. Carrier-sense for example, requires rapid feedback between the PHY and MAC in order to allow quick decisions to be made. And virtual-carrier sensing as mentioned above brings increases in power efficiency, but involves the PHY temporarily suspending operation in response to MAC commands. This requires tight interlinking, and is arguably the exact opposite of the total abstraction the layered model of networking purports to bring.

All that considered, while the breadth of terminology and variety of acronyms can make for some confusion; and indeed in the case of a large scale neighbourhood of devices, the amount actually going on can lead to further confusion; the actual distinguishing features of a wireless MAC protocol are few and simple, as has been demonstrated in this introduction to them. Furthermore, though practical implementation decisions in the majority of the communications devices currently in use mean that MACs and PHYs are thought of as inseparable pairs, genuine MAC/PHY inter-dependency and overlap necessary for functionality in legacy devices is usually only minor in scope. It will however be seen in the following section and in Chapter 3 that with the move towards CR, there are situations in which this overlap and inter-dependency grow much more necessary.

2.3 MAC for CR: What changes?

There is an exponential increase in flexibility and in possibility when considering the move from conventional MAC to MAC for CR. Whereas conventional MAC protocols for wireless communication are fixed by virtue of the hardware they are attached to, not to mention probable advance knowledge of the usage scenario, the enormous increase in the flexibility of a CR's physical hardware brings a commensurate increase in the flexibility of its MAC protocol. Even the use of the word protocol is perhaps no longer called for, as it implies something that is rigid and unchanging. The definition of what a frequency channel is; the length of time slot; the degree of time synchronisation, if any; the frequency of operation; the modulation order; the framing structure; the control mechanism; power; encoding scheme and many other parameters may all be subject to change in accordance with a CR's perceived needs.

Notwithstanding this increase in flexibility, just as the definition of a CR is open to interpretation as was discussed in §2.1.2, the definition of a CR MAC is also not consistent or commonly held. It is clear that its most fundamental change is that the medium itself can be redefined on-the-fly. But for a more concrete contrast with *legacy* MAC protocols for wireless communications discussed in the previous section, the changes in a CR MAC are best considered in light of the four basic medium access control issues presented in §2.2.1.

2.3.1 Redefining the MAC for CR: Reconsidering the basics

2.3.1.1 Choosing where and when to access the spectrum resource

The most fundamental change to the MAC occurs here. A CR MAC chooses not only where and when to access the spectrum resource, but *what* the spectrum resource *is*. Hypothetically speaking, a CR MAC may at any time employ any of the medium access techniques described in §2.2 in dividing the spectrum, and may change whenever it chooses. It is bound only by the capabilities of its PHY and the extent to which it can control those capabilities. Within these bounds, it may do as it wills.

Informing its choice is highly important, and therefore a CR MAC will be best served by having as great an awareness as possible of both its own operation and capability, and of its external environment. A CR must be aware of its RF spectral environment, and therefore spectrum sensing is an important responsibility. More than just a PHY layer concern, spectrum sensing requires the neighbourhood level picture available to the MAC layer, therefore a CR MAC will usually integrate some level of sensing coordination and sensing-based decision-making, as well as possibly sensing control. But its internal awareness is also of great importance. This is an awareness of the wealth of information possibly already available within a system, which can be used to make significantly improved decisions, and infer current and previous behaviour of other aspects of the system.

This information all goes into informing a CR MAC's choices, but perhaps the greatest challenge to be addressed is *how* it makes these choices. A CR MAC may employ advanced decision-making, machine-learning and Artificial Intelligence (AI) techniques to make its medium access decisions. These decisions may operate at different scales, some serving to choose the method of dividing up the medium, others to optimise individual parameters of this division. A knowledge base may be constructed in order to allow re-use of previous solutions that yielded good performance for similar scenarios. In this sense, informing decisions in a CR MAC may involve more than just using *current* awareness of internal behaviour and capability and current awareness of the environment and other possible devices within it; it involves awareness of *past* events, external and internal, considered not just in terms of what happened, but why it happened. It must be determined whether previous behaviour had the desired effect, and hence whether it is desirable to make use of it again. And on a further level, no decision-making mechanism is without its limitations and its costs. "Better" decisions are likely to take more time, and more computational resources to make. Considered further, it is up to a CR MAC to decide what constitutes a better decision for it, balancing its capabilities, the policies it must comply with, and the levels of performance it attempts to provide. Depending on circumstances, a rapidly made decision that makes reasonable use of the spectrum resource may be superior to one that requires significantly more time, but delivers excellent use of the resource. The best may not be better!

2.3.1.2 Ensuring message delivery

Ensuring message delivery is part and parcel of the choice of what spectrum resource to access. A CR MAC must make extra-efforts to ensure messages are delivered in an environment where future access to the current spectrum resource cannot be guaranteed. If it uses an ARQ-like mechanism, it must be capable of elegantly tolerating faults and disruptions, and resuming operation as quickly as possible.

Addressing also may become more complex, as in order to direct communications to their intended target, knowledge of the current status and spectral whereabouts of the target will be necessary.

2.3.1.3 Arranging necessary control signalling

Properly dealing with the above mentioned changes will require an increase in control signalling complexity, so that information is shared between devices in order to best facilitate neighbourhood awareness for decision-making. A CR MAC must balance the necessity of inter-device information-sharing with the overhead it involves. Furthermore, the MACs must arrange the means by which control signalling occurs, just as in a multi-channel MAC. The mechanisms introduced in §2.2.3 or similar ones will need to be employed.

In a greater sense, this particular responsibility expands to encompass not just control, but coordination. Such flexibility in medium access, and in the method of control signalling, means that devices will need explicitly arranged means to coordinate with one another in the event of neighbourhood changes. Mechanisms to enable network formation at device startup, as well as network recovery in the event of disruption, will be necessary. A trivial solution such as the assumption of the presence of a channel of guaranteed availability to use in order to initialise communication, may simply not be tenable[6]. Therefore more resilient, flexible and scenario independent approaches are likely to be required.

2.3.1.4 Handling the edge-cases of the chosen access mechanism

Both the variety and intensity of edge-case problems increase for CR MACs. Depending on the scenario, and the co-existence requirements, hidden and exposed node problems become more nuanced. This can be readily seen in the issues surrounding incumbent or primary-user protection in an overlay opportunistic DSA scenario. For example, if a CSMA/CA MAC is employed, a hidden node becomes a much more serious – and possibly illegal – violation of normal operations when it is a Primary User (PU) being interfered with. It could result in an active transmission being damaged by the CR system, and unbeknownst to it. Edge cases introduced by the coexistence of heterogeneous technology families, not just co-frequency but in adjacent frequencies, should also be taken into account. TDD and FDD systems, for example, may not naturally be suited to being adjacent in frequency.

2.3.2 Towards a concrete definition

While existing wireless MAC protocols are designed with specific scenarios in mind, a CR justifies its flexibility by being able to function well in a multitude of scenarios, and so its MAC must be more flexible. Examples of this scenario specificity could be: an assumption as to the maximum number of neighbours; the availability of certain frequencies for communication; the permission to transmit at some minimum arbitrary power level, etc.

There is also a greater blurring of the separation between the MAC and the PHY in a CR, as the MAC may wish to frequently demand parametric or even wholesale reconfigurations of the physical radio device. This crossover and blurring is a large part of what separates a CR MAC from a legacy MAC. A high level of control of

---

[6] Notwithstanding this currently unrealistic assumption, the appeal of such a default CCC as an enabler for CR is such that there have been tentative steps made towards the possible standardisation of one [28].

the PHY – which is admittedly also viewable as an interdependency – and a high level of PHY exposure to the MAC is what makes it capable of effecting its decisions, and of informing them. Every decision possibility, discussed in this chapter or hinted at, that in some way changes a spectral characteristic of the transmissions being made or prescribes reception behaviour – be it bandwidth, frequency, modulation order, modulation scheme, use of multiple carriers, decisions and methods of sensing, power control or many others – all will necessitate this kind of interdependency.

Returning to the confusion as to the definition of CR – and of CR MAC – it should be noted that all of the above is discussed in the context of this thesis's definition of CR. It is difficult at this point to proceed any further without an in-depth consideration of the literature, and in particular to examine the definition(s) of CR MAC held by its contributors.

An important question this thesis will attempt to address is the possible existence of an ideal CR MAC. That is to say, is scenario specificity acceptable? Can a CR simply possess a library of MAC protocols and choose the one most suitable to the situation it finds itself in based on past experience? Or is there one overarching super-protocol, which can be modified according to need and scenario, but is for the most part consistent, and best addresses most needs?

With this question in mind, the next section provides an overview of the state-of-the-art in CR MACs, with particular consideration of the employed definition of CR and CR MAC, as well as the common thread(s) between contributions.

2.4 CR MAC: The state of the art

With Cognitive Radio, Medium Access Control and MAC for CR discussed in the preceding sections, this section will present an overview of the current state of the art.

The literature concerning CR MAC is a significant body of work, and hundreds of papers have been published in recent years in the area[7]. A great many of these propose new CR MAC protocols, and many more offer solutions or mechanisms to specific problems considered by authors within the scope of CR MAC. The extent of this body of literature requires that any review be narrowed in its scope.

---

[7] A keyword search for the term *cognitive radio mac* currently returns in the order of 150 publications of the ACM and 400 of the IEEE.

The interest of this thesis lies in the definitions held by authors proposing works related to CR MAC, in the structure of their proposals, in what they consider the responsibilities of a CR MAC, and in the potential relationships between proposals in the literature. The details of specific CR MAC focused optimisation algorithms, or of niche analytical frameworks, are not of particular relevance to it. As a result, this literature review will focus on a cross-section of works self-declared as CR MAC. These have been selected where possible for innovative thinking in their definitions of the problems a CR MAC must face, and/ or of the solutions applied to these problems. Notwithstanding that, additional works deemed to fall under the same umbrella, but designated by their authors as falling under a narrower categorisation will also be discussed. The inconsistency in authors' definitions is no grounds for exclusion of novel and thought-provoking solutions, simply on the basis of differing or even excessively accurate labelling.

2.4.1 A selection of CR MAC protocols

27 protocols matching the above criteria were selected and examined and will be discussed in the remainder of this section. Rather than offer a per-protocol breakdown of specific details, it may be more instructive to examine their broadly common features and tendencies. This will lead to an overall characterisation of efforts to solve the challenges of CR MAC. In particular, researchers' definition of the role of a CR and a CR MAC, the flexibility of the protocol(s) they propose, the assumptions they commonly make, and the extent to which they investigate and implement their protocol(s) are among the prime factors of interest.

The tables overleaf offer a summary of the broad common tendencies of CR MAC research. For each complete protocol reviewed, Table 2.1 shows its definition of CR, definition of CR MAC, whether the protocol is actually opportunistic unlicensed overlay DSA, and the method(s) of evaluation employed.

It was alluded to in §2.1.2 that confusion as to the difference between CR and DSA is common. This is very much manifested in the literature, and it can be clearly seen from the table that research labelled as MAC for CR is in fact almost exclusively MAC for unlicensed opportunistic DSA. This is in spite of flexible definitions of CR sometimes provided. The only slight exceptions are in 3 of the 27 examined protocols. The protocol proposed by Felegyhazi *et al.* does not consider the existence of any primary users, instead discussing spectrum resource use between CR devices alone [29]. The protocol of Mišić *et al.* operates a *licensed* DSA scenario, in which primary users sell access to their spectrum, which is a sufficient variation from unlicensed DSA to be worthy of note [30]. And Cordeiro *et al.* present a protocol capable of functioning in a DSA scenario, but not reliant on it, and of sufficient flexibility to operate outside it [31].

In terms of the extent of their investigation, only 3 of the 27 examined protocols were implemented and tested using real-world radio equipment, even in a limited form. For the most part, evaluation was limited to simulation. While simulation offers a powerful insight into protocol performance, strict comparability is difficult to attain, especially when assumptions vary between authors. Furthermore, the examined protocols were not consistent in fully detailing the tools used for simulation, the parameters considered, and the full effects of the scenario assumptions being made. The effects and risks of such variety and over-reliance on simulation have been well investigated in the field of Mobile Ad-Hoc NETwork (MANET) research [32], and are equally applicable here. And while small-scale implementation as can be reasonably performed at a proof-of-concept stage suffers from significant challenges of its own – not least scalability, expense and experimental platform limitations – it adds significantly to the case for adoption of any protocol. Without it, there is a distinct risk that assumptions made in a protocol's design will go untested, and that its real-world usability may never be verified.

Table 2.2 notes some common simplifying assumptions made by many protocols, and some common features.

This second table clearly shows that all but a handful of protocols decide when, where and what to make of the spectrum resource by assuming frequency has been divided into orthogonal channels of equal bandwidth. Furthermore, the greatly simplifying assumption that all devices can hear the signals from all others – called a *Single Collision Domain (SCD)* – is made by 5 of the 27 protocols. Both of these are unrealistic and weaken the claims and applicability of the protocol, with the latter being particularly strong assumption, and highly unrepresentative of a real-world scenario.

Examining the common features noted in the second table, it is clear that the assurance of delivery by means of ARQ is common to almost every single protocol. To arrange control signaling [8], it can be seen from the table that virtually all proposed protocols require either a CCC, or slot time and use a dedicated portion of that time for common control communications – a Split-Phase (SP) control mech-

---

[8] This is marked "CSig" on the table owing to space constraints.

	Definition of CR			Definition of CR MAC					Is	Means of evaluation		
	a	b	c	i	ii	iii	iv	v	DSA	A	S	I
[33]	✓				✓				✓		✓	
[34]	✓				✓				✓	✓	✓	
[35]		✓						✓	✓	✓	✓	
[36]	✓			✓					✓	✓	✓	
[37]	✓			✓					✓		✓	
[38]	✓			✓					✓	✓	✓	
[39]			✓					✓	✓		✓	
[40]		✓		✓					✓		✓	
[41]	✓					✓			✓		✓	
[42]			✓					✓	✓		✓	✓
[43]	✓			✓					✓	✓	✓	
[44]	✓			✓					✓	✓	✓	
[45]	✓			✓					✓	✓	✓	
[46]			✓					✓	✓		✓	
[29]		✓					✓			✓	✓	
[47]	✓							✓	✓		✓	
[48]	✓							✓	✓		✓	
[30]	✓				✓				✓	In [49]	✓	
[50]	✓					✓			✓			✓
[51]			✓					✓	✓	✓	✓	
[52]			✓					✓	✓		✓	
[53]	✓							✓	✓	✓	✓	
[54]	✓			✓						✓	✓	
[31]		✓					✓			✓	✓	✓
[55]	✓			✓					✓	✓	✓	
[56]	✓							✓	✓		✓	
[57]		✓		✓					✓	✓	✓	

Definition of CR		Means of evaluation	
a	Opportunistic Overlay DSA	A	Analysis
b	A broader definition	S	Simulation
c	No definition provided	I	Implementation

Definition of CR MAC	
i	Access without interfering with PU(s)
ii	Choosing frequency to sense before access
iii	Extension of multi-channel MAC to avoid PU(s)
iv	A broader definition
v	Not provided

Table 2.1: Table detailing definitions of Cognitive Radio, definitions of CR MAC employed in referenced papers, noting whether protocol is unlicensed opportunistic overlay DSA, and noting method of protocol evaluation.

	Simplifying assumptions		Common approaches		
	Channelised	SCD	Uses ARQ	CSig method	Handles ECs
[33]	✓		n/c	n/c	
[34]	✓		✓	Hopping	✓
[35]	✓		n/c	CCC	
[36]	✓		✓	CCC or SP	✓
[37]	✓		✓	CCC & SP	✓
[38]	✓		✓	CCC	
[39]			✓	CCC	✓
[40]			✓	CCC	✓
[41]	✓			CCC	✓
[42]	✓		✓	Own model	✓
[43]	✓	✓	✓	CCC	✓
[44]	✓	✓	✓	CCC	
[45]	✓	✓	✓	CCC	
[46]	✓		✓	SP	✓
[29]		✓	n/a	Not needed	
[47]	✓	✓	n/c	CCC	
[48]	✓		✓	SP	✓
[30]	✓		n/a	SP	
[50]	✓		✓	CCC	✓
[51]	✓		n/c	CCC	✓
[52]	✓		n/c	CCC	✓
[53]	✓		✓	SP	✓
[54]	✓		✓	SP	
[31]	✓		✓	SP (Extended)	✓
[55]	✓		n/c	In-band	✓
[56]	✓		✓	CCC in ISM	✓
[57]	✓		n/c	CCC	

Definitions used in table	
Channelised	Assume frequency divided in advance into fixed, orthogonal blocks
SCD	Assume a Single Collision Domain
CSig	Control Signaling (See §2.2.1.3)
EC	Edge-case (See §2.2.1.4)
n/a	Not applicable
n/c	Not covered or discussed in the appropriate paper
CCC	Common Control Channel (See §2.2.3.2)
SP	Split-Phase (See §2.2.3.2)

Table 2.2: Table noting the prevalence of simplifying assumptions, and common approaches and features in the reviewed protocols.

anism – which therefore requires time sychronisation. Finally, the table shows that a great many protocols fail to address edge-cases at all. It is perhaps surprising that there is such uniformity of adoption of techniques devised originally for multi-channel MACs, and that only a small number of protocols employ novel methods for control signaling.

Overall, novelty for many protocols is found primarily in extended signalling mechanisms for handling PU activity [36, 37, 41, 50, 56]. For many others the novelty lies in specific mechanisms for choosing when to sense PUs [33, 34, 38]. But in summary, the limited degree of novelty manifested does not escape the static nature of these protocols, and does little to represent what might be considered a true CR MAC, of the kind that was discussed earlier in the chapter, in §2.3.2.

Perhaps the most salient overall observation that can be made echoes that made by Yau *et al.* [21] and inferred by Xiang *et al.* [58], namely that assumptions are for the most part so prevalent that the complexity of the problem is reduced to that of an advanced variant of a multi-channel MAC, possibly empowered with extra mechanisms to detect and protect primary users. It would seem that what is proposed in most of the literature is rather a scenario bound almost-CR, using "just another protocol", but still fundamentally incapable of reacting in a manner unforeseen by its designer(s).

2.4.2 Surveys & Tutorials

With a representative sample of CR MAC protocols from the literature discussed in the previous section, this section will discuss the observations made in surveys and tutorials covering CR MAC, a multitude of which have been published in recent years. These are works which discuss the broad expanse of work in CR MAC, usually providing a grounding context, before discussing and comparing the research of a variety of authors. It may be instructive to consider any overall observations made in these surveys, and in particular, it may be of interest to consider these surveys in terms of insights they offer in how they define CR, the challenges of CR MAC, and any methods of categorisation they use.

Several well-cited surveys do little to offer any insight into CR MAC beyond a broad background such as the one presented earlier in this chapter [9, 58–61]. These surveys do not provide novel categorisation schemes or address the challenge of what a CR MAC might be. They also almost universally equate CR with DSA.

Yau *et al.* observe the common challenges shared by DSA MACs and multi-channel MACs, though they refer to them interchangeably as CR MACs [21]. The authors argue that there are few effective differences between CR MACs, and multi-channel MACs, as they define them. This is an observation borne out and discussed in the previous section.

De Domenico *et al.* acknowledge the separation of protocols as currently presented in the literature into two classes, which they term direct access based and dynamic spectrum allocation based protocols. They do however remain consistent with other surveys in conflating CR MACs with DSA MACs, and furthermore they do not discuss the realism of the considered protocols, nor the possibility of implementation. Additionally, there is some confusion as to the separation of the classes into which they divide protocols. For example, the direct access based protocol class contains dynamic spectrum allocation based protocols. However the nature of the containment nor the distinction between protocols possessing this property is not clearly dealt with. Nonetheless this is an interesting observation, not made elsewhere.

The surveys & tutorials to be found in the literature do not therefore offer more than a limited insight over and above the discussion of the protocols discussed in the previous section.

2.4.3 Implementation of flexible MACs

Notwithstanding the observations of the preceding sections, there is a limited selection of research that focuses not on specific protocols, but on enabling realisation of flexible MAC protocols. This section will discuss this research and how it might provide a more fruitful path for exploration of the challenges of CR MAC. Of the ten papers discussed, only three are framed within the area of CR. Nonetheless, the remaining papers are very much of interest. In particular, most pre-date common knowledge of the CR paradigm and as such cannot be excluded on the basis of nomenclature.

TRUMP: Supporting Efficient Realisation of protocols for Cognitive Radio Networks [62]

In the first of the three papers framed within the context of CR, Zhang *et al.* propose and implement a toolchain to construct a CR MAC, that they dub "TRUMP". Using a simple C-like language, MAC state-machine elements can be arranged to form a

MAC protocol, which makes use of Rice University Wireless Open-Access Research Platform (WARP) hardware [63]. The protocol is then instantiated using a runtime engine on a combination of the WARP board's PowerPC core(s) and Field Programmable Gate Array (FPGA) fabric. The framework allows for conditional changes in the pathway to be followed, meaning that a TRUMP constructed MAC can alter the path of state-machine elements followed by a given packet, providing extensive protocol functionality flexibility. For example, a CSMA MAC can be altered to form a TDMA MAC by steering a packet through a *Schedule* function rather than a *CarrierSense* function. As in both cases the output of the function is a transmit decision made at an appropriate time, the remainder of the MAC's function is agnostic to and unaffected by the change.

This mechanism represents a particularly interesting proposal for its reduction of the complexity involved in making use of high-performance hardware. It effectively automates the process of implementing CR MAC protocols using the FPGA-based WARP platform, once they have been designed and transformed to suit the available state-machine elements. Provided they can be realised from the constituent blocks available, TRUMP allows for the rapid implementation of potentially *flexible* CR MAC protocols. Fig. 2.6 depicts the TRUMP development process, reproduced from the paper proposing it [62]. It shows the extent to which the framework allows easy access to hardware resources, given only a software protocol description. FPGA development is normally associated with a significant learning-curve, and therefore any such tools capable of reducing it significantly aid in experimental implementation. It should be noted however that in its current form, there are limits to the level of flexibility available from the TRUMP toolchain, and that in particular though it offers good timing characteristics, it restricts the level of control of PHY layer features, largely treating it as a black-box element, meaning that it is fixed and inflexible. In particular, it operates at only a fixed frequency block-size. This may be too stringent a limitation to implement some CR MAC protocols.

MultiMAC – An Adaptive Mac Framework for Dynamic Radio Networking [64]

In the second of three papers targeted at CR MAC, Doerr *et al.* firmly define a CR MAC as responsible for efficient reconfiguration using computational intelligence to make smart decisions about which of a suite of complete MAC protocols should be used, as well as which PHY properties should be set. As such they clearly agree with the "library of protocols" view of CR MAC discussed earlier, in §2.3.2. They propose an adaptive MAC framework for dynamic radio networking which is intended to evolve towards true CR Medium Access Control. Their "MultiMac" can dynam-

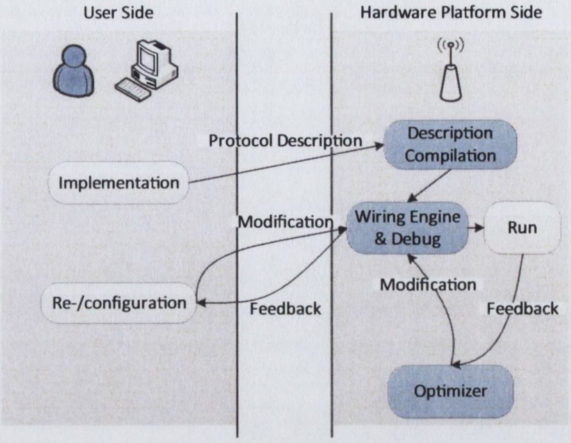

Figure 2.6: The development process of the TRUMP toolchain for MAC implementation, reproduced from the paper proposing it [62].

ically reconfigure MAC and PHY layers on a per-packet basis in response to local metrics and events. Effectively encapsulating many MAC protocols, the decision of which protocol to deploy is transparent to higher layers, and is performed by inbuilt optimisation mechanisms, policies for which can be set as desired by the user. The authors present an initial implementation using modified IEEE802.11 system.

Though defined by its authors as specifically within the context of "dynamic radio network", this proposal or meta-protocol and this thought process is of particular interest to the longer term evolution of CR MAC. Efforts to implement the system also assist substantially in demonstrating that it is a real-world possibility.

MAC Protocol Adaptation in CR Networks: An Experimental Study [65]

In the third of the three CR MAC targeted papers, Huang *et al.* also extoll a system based on a "bank of parametrically adjustable MACs". They dub the end result an "adaptive MAC protocol" which exists on a network-wide "global control plane." This control plane was introduced in previous work by these authors, and serves to handle all control and coordination across a CR network, not just at the MAC [66]. It is allocated an independent control channel. This control plane actively reconfigures the MAC protocol and PHY parameters of the data plane if a performance threshold is breached. The control plane is network wide, and a sufficient number of connected devices must vote for the change in MAC protocol – a localised drop in performance is insufficient – in order for it to be actioned. Like Doerr *et al.* with their MultiMAC [64], the authors of this proposal also discuss implementation, making use of a GNURadio platform. The protocol is interesting in that the control plane addresses a number of the challenges relating to network setup and consistent communication, as well as how to choose and be aware of MAC protocol choices across groups of devices. Also noted by the authors, and of interest in general is the possibility of propagating cross-layer statistics via the control plane, which can be used to make MAC protocol choice decisions taking into account a complex range of metrics.

The Click Modular Router [67] & Early Results on Hydra: A Flexible MAC/PHY Multihop Testbed [68]

Proposed in a different field, but of note as it offers some insight into the implementation and testing of communications protocols, is the popular "Click" modular router framework proposed by Kohler *et al.* Designed primarily for network layer implementation, it employs a modular approach to protocol construction, in which routing protocols are constructed from a library of blocks. In the context of wireless experimentation, it has been employed in tandem with GNU Radio by Mandke *et al.* The end result is a system that is similar to TRUMP, as discussed earlier, in that it allows for relatively rapid real-world implementation of test protocols. Flexibility and run-time change is not built into the system to the same extent as with TRUMP however, and so it is unclear the extent to which this may be of use in the context of CR MAC. Also notable is that GNU Radio's stream-based data-passing model does not necessarily suit it to the packet-based processing model of a MAC protocol. Nonetheless, such a modular approach to protocol implementation may well be a candidate platform for the enabling of CR MAC.

Meta-MAC protocols: Automatic Combination of Mac Protocols to Optimise Performance for Unknown Conditions [69]

Faragó *et al.* in their proposed "MetaMAC" acknowledge that a single protocol is unlikely to be best suited to all situations. In a proposal predating widespread adoption of CR, their solution is to operate multiple protocols simultaneously, combining the binary or probabilistic decision of each into a single decision based on the recorded successes of previous decisions, relevant to a single time-slot. It represents an interesting proposal, but suffers from a number of drawbacks. Time must be slotted for the mechanism to operate, which already represents a significant limitation of functionality. There is also no discussion of the computational realism of proposing that so many protocols be operating simultaneously. As it is a proposal that was made prior to the rise in prominence of CR as an idea, it was not strictly designed as CR MAC research. Therefore the protocols which are combined to result in an access decision appear very much self-contained, with little discussion or consideration of the variety of input these protocols might require to make their decisions. The system input in a CR MAC is an area of particular interest in that there may be such a range of them available.

Therefore it represents an interesting proposal, in particular in respect of the idea of a medium access decision being a built of a consensus based on experience. This is the sort of thinking that might inform the operation of an ideal CR MAC, but the method proposed in this particular work would require significant modification.

Automated Development of Cooperative MAC Protocols [70]

Lichte *et al.* propose a mechanism which could be the basis for the automated construction of MAC protocols, targeting cooperative MAC[9] specifically. They propose semantics to describe specific MAC handshakes, which can be daisy-chained together and then processed by a compiler to produce MAC state machines with complex overall behaviour. While thorough in its exploration, the specific methods proposed by the authors result in a mechanism for the control of a single-channel static PHY-layer in a CSMA-like manner. This is quite limited. In general however, the authors' argument is that the complexity of Cooperative MACs means that the direct creation of every individual handshake adds greatly to implementation complexity. They argue that a means to describe large swathes of intricate behaviour is highly necessary. This same argument could be applied to CR MAC, and therefore their primary contribution in the context of this thesis is this point.

Generalized Tree Multiple Access Protocols in Packet Switching Networks [71]

Sun *et al.* generalise MAC through the use of tree expansion, in which medium access decisions are split into multiple levels potentially corresponding to every way in which the medium can be divided. They divide possible medium access situations into either collision anticipation or collision resolution trees, with devices being grouped into leaves depending on the outcome of the algorithms of the chosen protocol. This provides an intriguing mechanism with which to analyse overall MAC behaviour, and to describe the group effect of a protocol. It does not accurately describe per-device protocol behaviour, but rather the emergent behaviour of the system as a whole.

A Unified Algorithm for Wireless MAC Protocols [72]

Teng *et al.* build on this work and use it to describe a protocol in terms of repeating

---

[9] Cooperative MAC involves the integration of retransmissions which are strategically designed to take advantage of possible diversity gains into the MAC protocol. Such protocols represent a perfect example of a cross-layer MAC-PHY construct, which brings benefit to both.

cycles of actions, and build a unified protocol capable of behaving like ALOHA, CSMA or Group Randomly Addressed Polling (GRAP). The novelty of the designed algorithms are limited by their effectively centralised description, however, and their restriction to a slotted environment. Though similar to the proposal by Lichte *et al.* discussed above in this respect, this proposal relates to that of Sun *et al.* in that it also focuses on network-wide emergent behaviour.

The proposals of both Teng *et al.* and Sun *et al.* predate the common acceptance of the CR concept. While they might represent an interesting network-wide model for analysis, it is difficult to see how either would be directly applied to the implementation or design of CR MACs.

A State-Machine Based Design of Adaptive Wireless MAC Layer [73]

Xiao *et al.* propose a state-machine based framework, pointing out seven common states which can be parameterised and used to construct a range of MAC protocols. Similar to others, the authors here argue that reconfiguration at the MAC layer is the best solution for handling heterogenous capability. Beyond the assumption of a rapid MAC to PHY to MAC turnaround however, there is little in the way of PHY interaction. The range of MACs it is possible to implement using this method is limited to those entirely abstracted from specific PHY requirements. The authors provide an example discussion of an implementation of a General Packet Radio Service (GPRS) MAC using their method, a centralised multi-channel MAC. In general however, the method does not seem suited to deployment for distributed multi-channel MACs, and the PHY interactions involved in a multi-channel MAC are not discussed. Nonetheless it is an interesting mechanism, similar in its thinking to TRUMP, discussed above.

These protocol papers provide specific elements that represent innovative solutions to DSA and CR MAC problems, and represent innovative and complete thinking as to the mechanisms that might offer true flexibility from circumstance to circumstance. This contrasts strongly with the scenario-bound protocols discussed earlier, in §2.4.1, which are fundamentally incapable of reacting in a manner unforeseen by their designer(s).

2.5 Discussion

In consideration of the introductions earlier in the chapter, and of the literature review, several key points come to light.

Any MAC protocol for wireless has relatively few options in terms of how it realises medium access. Protocols may emerge as complex and detailed, but ultimately represent compound bodies of simple elements – ordering symbols and the bits they represent in a manner chosen to effect consistent and largely interference free or interference- limited access to the medium.

CR MACs can be significantly more complex, as they are part of a device in which the meaning, characteristics and context of symbols is subject to change in order to suit circumstances, and radio features may be employed to gather information not available to a legacy MAC.

It is difficult to consider in strict terms the definition of a CR MAC. The short answer may be "a module or mechanism that minimises inter-device interference", but really it is a mechanism that in a distributed local-neighbourhood sense changes as and when it sees fit the mechanism for access while obeying externally set policies and physical capability restrictions. How to actually implement such a thing is an issue largely avoided.

The simple fact is that the vast majority of research situated as CR MAC research is in fact DSA, and for the most part overlay DSA at that. This is not to decry DSA targeted research. Nonetheless the lack of recognition of the separation perpetuates confusion as to the difference between DSA and CR, and sidesteps some of the larger research questions to be asked in the broader context of CR MAC. Furthermore, as first observed by Yau *et al.* [21], the similarities between most proposed CR/ DSA MACs and multi-channel MACs are uncanny. The assumptions made by protocol designers, and the reduction of scope resulting from them means that in many cases the ultimate contributions are of arguable novelty even for multi-channel MACs.

The literature has sidestepped the broader issues by framing itself in terms of such specific scenarios that contributions are largely incomparable. Each protocol is predicated on a series of assumptions, which its "competitor" protocols are unlikely to share. Though this is common practice with a great deal of research, it is nonetheless unhelpful in advancing the broader field. Additionally with underspecified and inconsistent simulation and analysis the primary methods of evaluation employed, the degree of incomparability increases further. This continues both as the model employed for evaluation varies for each proposed protocol, and as authors rarely compare their protocols with those of others. Simulation can easily be rendered questionable as a proof of concept, as is well detailed for MANET research [32]. However, even if it is well executed there is a strong need for increased implementation in order to mitigate real-world issues and challenges, in addition to rigorous analysis to assess algorithm deficiencies and capabilities.

Implementation and experimentation remain the ultimate test of any proposed system or method, and philosophically speaking are usually the stated end-goal of any invention. Implementation and experimentation in CR MAC remains limited however, and even experimentation that has been performed is for the most part employing inflexible Consumer Off-The-Shelf (COTS) hardware such as 802.11 hardware paired with the MadWifi flexible driver [74], custom GNURadio implementations or in limited instances, WARP hardware. Proposed methods of addressing some of the specific challenges in CR MAC implementation are introduced in [75–78], whose authors in particular note the difficulty of implementation on limited GPP attached SDR hardware.

The research discussed in §2.4.3 does represent the first steps towards providing a more useful platform for CR MAC implementation. More broadly, it opens the discussion as to how a CR MAC might look "under the hood", in terms of the range of protocol features that might change, and how the change might be effected. This small body of research seems very much divorced from the static protocols discussed in §2.4.1. Each one of the discussed static protocols is introduced fully formed, with no explanation of the process that led to its design and construction. By contrast, the research in §2.4.3 opens the conversation on processes involved in constructing CR MAC protocols. It is true to say that several of these proposals choose to deliver flexibility by encompassing and switching between several wholly contained protocols, and that construction of these internal protocols is itself not discussed. But proposals such as TRUMP [62] and the modification to Click suggested by Mandke *et al.* [68], provide a toolchain for construction of protocols from scratch. However, even these shy from the larger context of CR MAC, of what the constituent parts are being built to do. They provide the "ingredients" for a CR MAC, but do not discuss how to use them to create a recipe.

To return to our definition of it, in a CR as a whole, either a Cognitive Engine or some other body will be required to make decisions as to when and what reconfigurations occur. As discussed in §2.2.7, MACs for wireless communication already blur the boundary between MAC and PHY. As further mentioned in §2.3, CR MAC only serves to further denude this boundary. It is fair to say that the MAC may well be involved in that role of controlling reconfigurations, or possibly even possess it wholly. This will requires intimate radio knowledge, a capability for self query and an internal model not dissimilar to that originally proposed by Mitola [1], and discussed by Lotze [79]. Its secondary effect is an increase in the understanding of the whole radio device by MAC designers.

In spite of the broad scope of discussion in this chapter, it is fair to say that enormous conceptual leaps are unlikely in this field. It is improbable that a CR MAC will apply radically *new* techniques for multiplexing – it will simply apply existing ones, changing them over time if it deems that prudent. Exciting advances such as those referred to in §2.2.6 may expand the range of capabilities available to a CR MAC slightly further, but true research advances will come from a CR MAC's interaction with the cognitive radio as a whole, and in particular how it best harnesses its capabilities.

2.6 Therefore

Therefore the aim of this thesis is to explore MAC for CR, with a view to expanding the understandability and usability of of existing CR MAC research and also to aiding the path of future research. In particular, it aims to assist with the implementability of existing and future research. It aims to unshackle protocols from the specific scenarios envisioned by their authors, and to explore the steps necessary to enable movement towards a MAC for a CR that closer meets the ideal. The next chapter will begin by introducing a number concerns which stand in the way of these these aims, and then propose a mechanism to address them.

3 The Anatomy of a Cognitive Radio MAC

3.1 Making more of CR MAC research

Subsequent to its introduction of Cognitive Radio (CR) Medium Access Control (MAC), the previous chapter discussed issues presenting themselves in and around the vast majority of CR MAC protocol proposals in the literature, and works discussing the challenges of CR and flexible MAC realisation in general. These issues included the difficulty of comparing protocols; the relative lack of novelty of protocols; a great variety in grounding assumptions harming both comparability and usefulness of protocols; and overarching them all a large degree of disagreement as to what a CR – and hence a CR MAC – actually is. This chapter aims to maximise the usability of CR MAC research thus far, and also seeks to increase the potential benefit of future research. It will do so by advancing in detail a number of core concerns that must be addressed to facilitate these aims, and then introduce and explain a mechanism which does so – the *Anatomy of a CR MAC*.

3.1.1 Concern: The case for deconstruction and reconstruction

The first concern to be discussed is the method of presentation of most of the proposed CR MAC protocols. They are described as monolithic entities. Usually, novelty lies only in one or two small parts of the overall protocol, the remainder of which offers little to no contribution. Given that, it would be highly desirable to be able to extract the novel portion from the overall protocol, both to consider it in isolation, and also to consider how it might be included as part of an alternative protocol. Rather than being bound by the arbitrary choices of authors, in terms of what protocol they build around the items of novelty, it should be left open to include these novel portions in *any* protocol they can safely function as part of. The process of division of a protocol into parts, and the consideration of those parts in isolation shall be termed *deconstruction*.

Given a consistent and sufficiently well defined method for deconstruction, these protocol portions could be examined individually, and integrated into existing or new CR MAC protocols. They would be capable of performing the same role in potentially very different settings. And provided the purpose of a portion is clear, functionality of a MAC protocol would be comparable on a much more meaningful level than currently, where the whole protocol must be considered as is, and any attempt to deconstruct it is haphazard.

3.1.1.1 Dependencies

If a CR MAC protocol is divided into smaller portions, it is misleading to say that they can stand completely alone. To consider the picture more fully, a CR MAC *can* accurately be divided into constituent portions, but these portions will unsurprisingly be interconnected. They are *dependent*. If there is to be a consistent and well-defined method for deconstruction, this method must also take into account and make clear the interconnectedness of CR MAC protocol parts, and take into account their dependencies. When comparing between CR MAC portions from different protocols, the parts' respective dependencies must also be compared, in order to ensure that reasonable and accurate parallels are being drawn.

3.1.1.2 The possibility of reconstruction

Provided deconstruction has taken place with sufficient consistency and with awareness of dependencies, it is not unreasonable to presume that a new possibility opens up – that is, the possibility of *reconstruction*. Reconstruction is the combination of CR MAC parts that resulted from deconstruction. A group of such parts might be combined to create an entirely new protocol from the ground up, or a part may be inserted into an existing CR MAC protocol, in order to expand its functionality.

This possibility of reconstruction would greatly increase the ease with which CR MACs can be constructed, allowing them to be built from a range of novel protocol parts. These parts might originally each have been proposed in the context of its own protocol. Irrespective of even strong differences between these source protocols, it may be possible to combine them. As importantly, authors of CR MAC protocols whose monolithic nature may previously have "trapped" useful functionality might find deconstruction and reconstruction enabling an uptake in use of their research. Fully enabled, the end result would be to allow the novel portions of existing and future works of CR MAC research to be used in ways that need not be foreseen by their authors.

3.1.1.3 Dependency conflict

It is a trivial observation that not every CR MAC protocol part can be combined with any other. Therefore rules for combination must combat the risk of *dependency conflict*, in which the requirements of two components are mutually exclusive.

In summary, this thesis moves that a means to deconstruct CR MAC protocols into their constituent parts would contribute significantly to making existing and future CR MAC research more understandable, more usable, more implementable and more future-proof.

3.1.2 Concern: Spectrum Resource Allocation vs Exploitation

This chapter has already mentioned that a great many CR MAC protocols differ in their definition of what a CR is. What is further of import is that even when protocols agree as to what a CR is, they may propose CR MAC protocols with very different roles. In fact, there are two distinct roles that can be played.

3.1.2.1 Division of spectrum in legacy systems

Outside of Cognitive Radio, systems designers build MACs for a physical (PHY) layer that is already characterised, and for a radio frequency (RF) environment that is known. This is because interference is managed externally, usually by a government regulatory body. Technology operators are granted spectrum licenses by these regulators overwhelmingly on a command-and-control basis that has been employed for many decades. That is to say that blocks of frequency are granted for use to licensees subject to the use of a specific technology and with exacting power control requirements and out-of-band emission limitations. These licenses are non-transferable in any timeframe a MAC may ever be reasonably expected to deal with, and so represent fixed specifications to be incorporated into a system.

This command-and-control style of thinking extends right into the design of legacy wireless systems. Consider an infrastructure network of wireless devices, in which Access Points (APs) or base-stations[1] communicate with connected devices. Spectrum resources must be divided among APs in a manner that prevents an AP from interfering with its neighbours. Once these resources have been appropriately apportioned to an AP, it can then in turn divide them among its connected client devices. In legacy systems this division among APs is performed in a static manner. For example, in cellular networks such as GSM or UMTS the division of resources between APs is manually planned by the network's designers, and usually remains fixed. And in IEEE802.11 devices, the desired channel on which an AP operates is selected by the user, and then remains static during ordinary operation.

In a CR device, the fixed and highly controlled environment that legacy wireless devices assume, is no longer present. Therefore division among APs can no longer be static, and must become the responsibility of the CR MAC.

3.1.2.2 Spectrum resource allocation: An additional role for CR MAC

If the infrastructure network discussed in the previous section were reconsidered as a network of CRs, the devices' MAC protocols would be required to play two roles.

The first is its traditional MAC role, which this thesis terms *spectrum resource exploitation*. To play this role, each AP divides the spectrum resource between client devices, to facilitate their communication with the AP and possibly one another.

The second is the new role described above, that of *spectrum resource allocation*. This is the process of dividing the spectrum among APs, so that they avoid excessively interfering with one another.

3.1.2.3 A CR MAC's two roles

When both roles of a CR MAC are appropriately played, RF spectrum resources can be allocated coarsely among APs, and then exploited by APs and clients as their respective MACs see fit.

These roles can also map to a distributed scenario. In such a scenario, spectrum resource allocation takes place in terms of transmitter-receiver links, or of groups of CR devices. Spectrum resource exploitation determines the use of the resource within these links or groups to communicate.

---

[1] The terms Access Point (AP) and base-station will be used interchangeably from now on.

	Exploitation	Allocation
[33]	✓	
[34]	✓	
[35]		✓
[36]	✓	
[37]	✓	
[38]	✓	
[39]	✓	
[40]		✓
[41]	✓	
[42]	✓	✓
[43]	✓	
[44]	✓	
[45]	✓	
[46]	✓	
[29]	✓	✓
[47]	✓	✓
[48]	✓	
[30]	✓	
[50]	✓	
[51]		✓
[52]	✓	✓
[53]	✓	
[54]	✓	
[31]	✓	✓
[55]		✓
[56]	✓	
[57]	✓	

Table 3.1: Table of complete protocols discussed in Chapter 2, §2.4, broken down in terms of CR MAC roles played.

With the distinction between these two roles clarified, the literature reviewed in §2.4 of the previous chapter can be reexamined in terms which role(s) each protocol plays. This information is shown in Table 3.1. Upon examination of it, it becomes clear that of the complete protocols discussed, some play a resource allocation role only, and others a solely resource exploitation role. Only 5 of the 27 play both.

Lacking in the literature is this realisation that ideally most if not all CR MACs should perform both roles. In fact, recognition of the existence of and distinction between these roles is relatively rare in the literature altogether. Proposed protocols are not explained or situated in terms of which role they play. Instead all are described generically as CR – or in some cases opportunistic Dynamic Spectrum Access (DSA) – MAC protocols.

It should be noted that in a very limited number of cases, these two roles are distinguished. A division of existing CR MAC protocols into similar categories is made by De Domenico *et al.* in their survey, described by the authors as Dynamic Spectrum Allocation and Direct Access Based MAC protocols [80]. Similar groupings using different but analogous terms are used to frame the protocol in [52]. But in such cases the roles are considered as a grounds for categorisation of self-contained CR MAC protocols. This thesis argues that in order to represent a complete CR MAC, and in particular come close to an ideal one, a system must play both a resource allocation and exploitation role. The former works to manage co-existence and the coarse-grained allocation of spectrum between groups of devices; the latter to decide how spectrum is divided among devices in the same group for communication.

Given that protocols playing these two roles are for the most part proposed separately in the literature, it is worth noting that in such a complete CR MAC – which as defined above, plays both roles – some potentially unresolved difficulties may arise. Firstly, as spectrum resource allocation occurs among groups of two or more devices, the makeup of the group must be decided. Secondly, only one device in each such group will need perform the allocation. Which group member it is, must somehow be determined. This choice is usually trivial in an AP-client or somewhat centralised scenario: the AP is likely to be the best allocator. But in a decentralised scenario, there is need for some distributed mechanism for choosing an allocator, and of replacing it in the event of a change of circumstances and/ or group membership. The final challenge in such a situation is how to disseminate the results of the allocation mechanism to the non-allocating devices in the group. Though there is little to no discussion of these challenges in the literature, a complete CR MAC must address them.

In summary, a mechanism to allow easy recognition of the role a protocol plays, and to facilitate development of protocols playing both roles, would greatly aid in making CR MAC research more understandable, usable, implementable and future-proof.

3.1.3 Concern: PHY layer dependencies

As alluded to in Chapter 2, §2.2.0.1 and §2.2.7; for cost, performance and efficiency reasons the MACs of virtually all RF-based communications devices currently in use are implemented on Application Specific Integrated Circuits (ASICs). This places them on the same silicon on which PHY signal-processing occurs. Furthermore, in many cases, close examination of the on-chip design will reveal that there is no strict separation between the MAC and PHY functions. In fact, processing elements are shared, and there will be a high degree of feedback and pipelining, all with the aim of improving performance, reducing power consumption and functional redundancy, and minimising the amount of silicon real-estate in use.

Though designed conceptually in isolation, these fixed and highly coupled implementations of legacy MACs mean that their designers are accustomed to having direct and high-speed access to PHY features, in addition to having the deterministic performance that dedicated hardware allows. When such MACs are considered in abstract terms, it is important to be aware of the implicit *physical layer dependencies* that this results in.

An example is the Carrier Sense Multiple Access with Collision Avoidance MAC, or CSMA/CA, as employed by the IEEE802.11 wireless networking standards [3]. When employing this protocol, it is crucial to remember that i.) the latency between deciding to transmit and the commencement of actual transmission must be minimal and ii.) the time taken for a received packet to reach and be processed by the MAC must also be minimal. Otherwise a device may sense the medium idle and request transmission of its PHY, but be delayed by hardware latency to the extent that another device can commence transmission. Such an occurrence would result in a high probability of collision for both devices. If this delay were consistent and known across all devices, the problem could be mitigated with an altered protocol algorithm. However, this issue of turnaround time is a particular problem for Software Defined Radios (SDRs) where signal processing takes place on a General Purpose Processor (GPP), but the RF front-end hardware is attached via a bus which suffers from a relatively large and non-deterministic delay [77]. While the size of the delay causes problems, greater impact comes from its unpredictability and inconsistency, as protocol algorithms cannot be easily or reliably altered to account for unpredictable delays. Furthermore, when this problem is scaled up to a network of such SDR-based devices using CSMA/CA, heterogeneous behaviour will lead to high performance-jitter even in controlled scenarios. This requirement for quick and consistent turnaround-time is a distinct physical layer dependency.

Another example of a common dependency embedded in a great many published CR MAC protocols is the reliance on synchronisation from a time slotted PHY (and/ or Primary User (PU)). This synchronisation is usually used to simplify decisions as to when and at what frequency spectrum occupancy should be observed [59, 60, 80].

Awareness of and frank discussion as to the level of PHY dependency of a protocol is not something found in the majority of works in the literature proposing CR MAC protocols, however. This is a serious deficit, as the potential level of PHY dependency increases enormously in a CR MAC. Cross-layer interactions in legacy MAC protocols may have been made with a view to performance improvements, or to address edge-cases. In a CR, in order to contribute fully to deciding what the spectrum resource is, as well as when and where to access it, a MAC protocol will likely require extensive control of PHY features, deep access to PHY metrics, and complex PHY performance and capability guarantees. The examples of synchronisation and rapid-turnaround provided above apply to CR MAC to at least the extent they do for legacy MAC protocols. But to reiterate the list of possible PHY features from Chapter 2, §2.3.2, a CR MAC protocol may also require access to control of and metrics relating to bandwidth, frequency, modulation order, modulation scheme, use of multiple carriers, decisions and methods of sensing, power control, to list but a few.

In order to make CR MAC research more understandable, usable and implementible an awareness of these dependencies is highly important. They are an integral part of a CR MAC protocol, and are necessary to allow accurate comparisons between protocols, in addition to being required to allow implementation [81].

When a CR MAC protocol is being considered in terms of its parts, as motivated above in §3.1.1, a consideration of PHY dependencies continues to be of great importance. In a deconstructed CR MAC, PHY dependencies should be considered in terms of the protocol parts that cause them, not as properties of the protocol as a whole. Each part contributes separately to the overall set of dependencies – some parts may have none, while another part may be the cause of many or all of them.

Full specification of a complete CR MAC protocol, and accurate deconstruction of a protocol require that PHY dependencies be taken into account. Therefore doing so is unquestionably a required step in making CR MAC research more understandable, more usable, more implementable and more future-proof.

3.1.4 Concern: Being aware of awareness

Finally, an enormously important aspect of CR is *awareness*. The very intention of a CR is to be a device aware of its surroundings, but also of itself, its capabilities and its effects on those surroundings both current and potential. This aspect of it more than any other is what separates a CR from a legacy device, even one capable of great flexibility. The concern that relates to a CR's awareness, is, paradoxically, the need to be aware of it. It is in fact an aspect that does not appear to be considered quite integral, and hence does not manifest itself in the literature as often as might be expected.

Only one area in particular of CR MAC awareness stands out in terms of its prevalence in the literature, and that is spectrum sensing. Most spectrum sensing relates specifically to the scenario of opportunistic DSA. Spectrum sensing in general can be considered an awareness of the world *external* to the CR device, as it relates to the current and/ or predicted status of other devices and/ or networks of devices, often non-CR. It is fair to say a great deal of research has been focused and indeed continues in this area [8], as it does in the alternative mechanism of spectrum occupancy databases. And while both are of particular importance to opportunistic DSA, the great many other mechanisms that may fall under the heading of awareness should not be forgotten, as they are potential sources of highly interesting and innovative behaviour and capability.

The entire range of possibilities relating to awareness *within* a device and its neighbouring CR devices, for example, offers exciting possibilities and information sources to inform complex and nuanced CR MAC decisions. Furthermore, it's important to note that awareness in this context relates to all resources available to a CR, not just the spectral. Internally, these resources might include computational power, leading a CR MAC to make use of simpler Digital Signal Processing (DSP) that might not yield the most efficient signal to be transmitted, but will not leave the medium unduly idle. They might include energy availability, leading a CR to use simpler and more energy efficient modulation techniques, or choose a frequency with better signal propagation in order to conserve battery-power. They might include knowledge of its own poor RF transmitter characteristics, leading a MAC to make extremely conservative access decisions in order to avoid damaging other devices and/ or technologies with its out-of-band emissions. These are but a few examples of the power an internal awareness imparts to a CR MAC. And of course there is external awareness beyond awareness of spectrum and spectrum resources. This could include position and topological awareness[2] which might allow use of line-of-sight communications techniques in wide-open environments, while preferring frequencies and modulation better suited to building penetration in urban areas. It could include awareness of temperature and its effects on the device's physical hardware. Or it might take into account human events, serving an altruistic relay role in the event of an emergency, or deferring from use of suddenly lifesaving radio resources.

In summary, the great concern relating to awareness in a CR, is the need for the CR to have it. It is an unavoidably integral part of a CR, and in particular a CR MAC and comprises significantly more than just spectrum sensing.

3.2 Addressing concerns: The Anatomy of a CR MAC

The previous section outlined a number of concerns that must be addressed in order to solve the problems with CR MAC research as it is currently conducted.

Fundamentally speaking, a lack of agreement as to what CR MAC entails is the root problem with current CR MAC research. Given that lack of agreement, the current method of presentation of contributions as complete protocols bound to a great variety of arbitrary scenarios buries their novelty, and frames contributions in very specific terms that make comparison and complete understanding difficult. This motivated the idea of deconstruction of CR MAC protocols. Furthermore, a CR MAC plays two roles, that of resource allocation and exploitation. Examples of MACs protocols performing one or the other are manifold in the literature. But no authors acknowledge that a CR MAC protocol should perform both. Finally, a CR MAC protocol very likely possesses multiple PHY dependencies which should be taken into account when evaluating, implementing and/ or deconstructing it.

A mechanism addressing these concerns would provide a mechanism for protocol deconstruction, in terms of the roles it plays, while also taking into account PHY interactions and dependencies. Therefore, this section proposes and explains such a mechanism, the *Anatomy of a Cognitive Radio MAC*.

---

[2] It should be particularly noted that any device which makes use of a spectrum occupancy database will likely have such information "for free", and therefore in many cases the availability of such awareness need not be considered onerous.

3.2.1 Introducing the Anatomy Schema

The Anatomy of a CR MAC is a schematic representation of the functionality of a CR MAC protocol, divided into its constituent parts, termed Anatomy *elements*. It details the communications data-pathways between these elements, and the outside-of-dataflow informational dependencies between elements. It divides these elements into the two roles a CR MAC protocol can play. And it clearly delineates the PHY dependencies of each element.

A template *Anatomy Schema* is shown in Fig. 3.1. Note that this example shows only one PHY transceiver, but a CR MAC might require multiple transceivers.

The remainder of this section will now discuss the features of the Anatomy Schema in detail, in terms of the concerns detailed earlier in the chapter.

3.2.1.1 Deconstruction & Reconstruction

As mentioned in the introduction to the Anatomy Schema above, it divides a CR MAC protocol into Anatomy elements. This division facilitates deconstruction as elements can be considered independently, provided their dependencies are sufficiently well defined. It also facilitates reconstruction, as groups of elements can be recomposed to create new protocols, or enhance existing ones, meaning the concerns of §3.1.1 are addressed.

The steps involved in deconstructing and reconstructing specific CR MAC protocol will be discussed later in this chapter in §3.3. It is important however, to consider the framework in which the Anatomy Schema places the elements of a CR MAC, which the remainder of this section will discuss.

Chapter 2, §2.2 introduced the four responsibilities of a MAC protocol for wireless as i.) choosing when and where to access the spectrum resource; ii.) ensuring message delivery; iii.) arranging necessary control signalling; and iv.) handling the edge-cases of the chosen access mechanism. These responsibilities also apply to CR MAC, with an expanded role for the first of them: In a CR, the MAC takes some responsibility for *defining what* the spectrum resource is.

This notion was introduced in Chapter 2, §2.3, but warrants further exploration. In legacy MAC for wireless devices, what the medium means is effectively defined by the PHY, in terms of bandwidth, modulation scheme, whether it is single or multi-carrier, etc. More advanced MAC protocols began to change this, with PHY

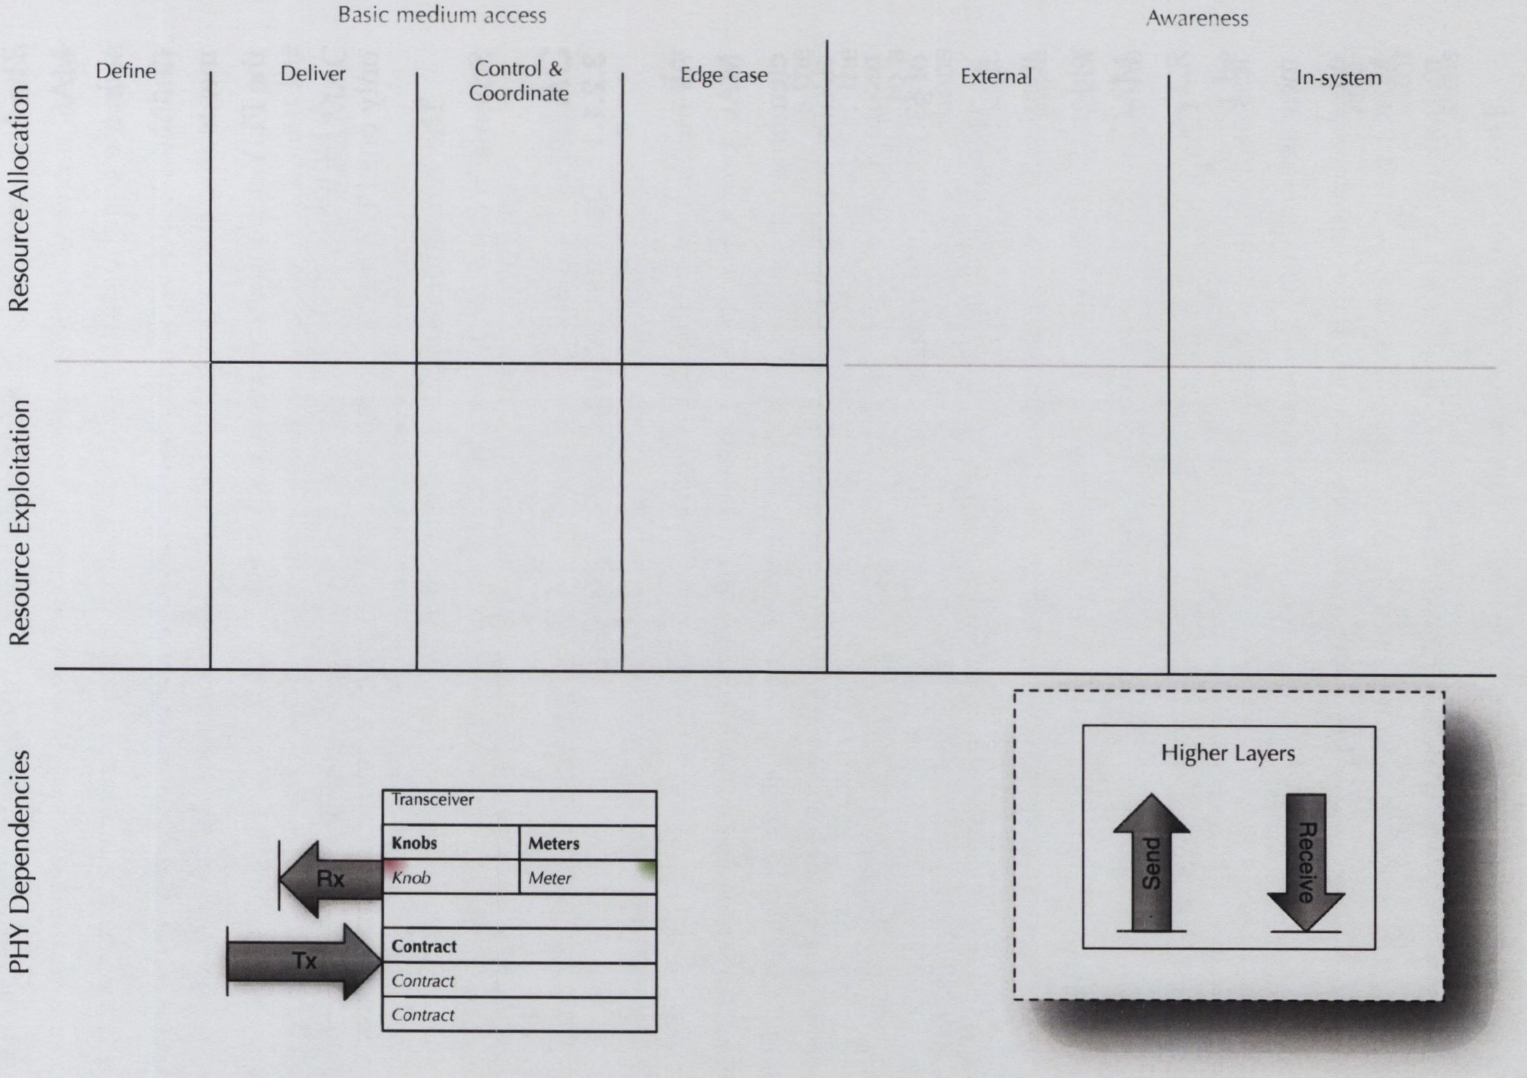

Figure 3.1: An uncompleted template Anatomy Schema for a single-transceiver CR MAC. Note the connection to higher layers is placed at bottom right to fit within available space. It is not in any way directly connected with the transceiver(s) to its left.

transceivers becoming more flexible. In IEEE802.11 for example, modulation order is determined by checking the Signal to Interference plus Noise Ratio (SINR). And in 802.11 revision 'N', bandwidth is also potentially unfixed, using 40MHz when communicating with other 802.11n devices, while using 20MHz in communication with legacy 802.11a/b/g devices [23].

CR will take this significantly further, potentially turning the paradigm on its head and leaving more of a communications specification unfixed than fixed. In this case, a CR MAC must also play a role in choosing PHY parameters, and hence defining what the spectrum resource is, as well as when and where to access it.

The remaining three responsibilities of a MAC also expand in their scope, as discussed in Chapter 2, §2.3. Delivery must be ensured in a much more unpredictable environment; control signalling expands to encompass coordination also, and requires devices to be able to discover and communicate with one another in a manner resilient to this unpredictable environment; edge-cases increase in their variety and in the risks they present.

As all of these responsibilities are important, the Anatomy of a CR MAC divides Anatomy elements under these four headings. They can be seen as four columns on the left hand side of the template Anatomy Schema shown in Fig. 3.1, namely:

- **Define**
- **Deliver**
- **Control & Coordinate**
- **Edge-case**

Each of these is short hand for its respective MAC responsibility. Many Anatomy elements into which a CR MAC is divided will usually fall under one of these responsibilities, and so can be arranged in the appropriate column on the template.

3.2.1.2 Resource Allocation & Exploitation

As introduced in §3.1.2, a CR MAC protocol may play two roles, and ideally should play both. These roles are that of resource allocation and exploitation. Therefore, in addition to considering Anatomy elements in terms of the four responsibilities of a CR MAC protocol, the Anatomy Schema also divides them in terms of the role they contribute to. This immediately clarifies which role(s) a protocol is playing.

Elements in an Anatomy Schema are laid out in a grid, and within both of the rows corresponding to the two roles, the responsibilities of defining, delivering, controlling/ coordinating, and handling edge-cases will ideally or at least mostly be met by suitable elements.

As an example, consider the protocol proposed by Cordeiro *et al.* in [31], and discussed in Chapter 2, §2.4. Note from Table 3.1 shown earlier, that it plays both CR MAC roles. Furthermore, it does so in a distributed scenario. In playing its allocation role, it defines frequency channels that are allocated to groups of devices. It facilitates control and coordination between these groups by means of a combination of a selected "rendezvous channel", and rules for beacon synchronisation. In playing its exploitation role, it defines when to access the spectrum resource for each device in a group by slotting time in its allocated channel. To ensure delivery, it employs Automatic Repeat reQuest (ARQ). For control and coordination a portion of the slotted time is dedicated to control communications, and it defines a number of edge-case mitigation mechanisms relating to recovery in the event of PU occupancy of a group's channel, and the possibility of groups becoming too large.

This example demonstrates the separation of elements corresponding to responsibilities between roles. That is, an element relating to a responsibility in a MAC's allocation role, does not necessarily have any involvement in that same responsibility in its exploitation role.

3.2.1.3 PHY dependencies and contracts

The Anatomy Schema acknowledges PHY dependencies, and furthermore acknowledges that they can be divided into two categories.

PHY transceiver(s) are marked on the Anatomy, beneath the rows corresponding to resource allocation and exploitation. Note that they do not bear any direct relationship with the columns corresponding to the four responsibilities of a MAC.

The first variety of PHY dependency which an Anatomy element may possess is behavioural. That is to say, an Anatomy element may rely on a consistent pattern of behaviour from a transceiver, or property of that transceiver, in order to function. The Anatomy terms such a behavioural requirement a *PHY contract*. Examples of a PHY contract might be the employment of a multi-carrier modulation scheme, the guarantee of inter-device time synchronisation or some performance guarantee such as latency or minimum bandwidth. Note that PHY contracts are passive. Once a transceiver has been contracted to behave in a prescribed manner, the behaviour requires no direct input from the CR MAC to sustain, instead it represents a constant behaviour that should be transparent to the MAC.

The second variety of dependency is interactive, and is itself divisible into two specific classes. The first class of interactive dependency is that of an Anatomy element on access to some state or measurement information at a PHY transceiver. The Anatomy terms this class of dependency as a *meter*, analogous to an old-fashioned measurement tool. Meters are a manner in which cross-layer awareness requirements can be represented. An example of a meter would be the requirement for access to estimated SINR information, the current signal bandwidth or some other variable system property of the transceiver. The second class of interactive dependency is that of an Anatomy element on the ability to cause change in a PHY transceiver. It is easy to see how such a dependency is common-place in a CR MAC, in light of its additional role defining the medium it controls access to. Analogous to the means of adjustment in an old-fashioned radio, the Anatomy terms this class of dependency a *knob*[3]. Examples of a knob would be an Anatomy element requiring the ability to change PHY frequency, or activate and de-activate Forward Error Correction (FEC) coding.

The second of these two offers a slightly more complex example, as for a transceiver to present the ability to activate and deactivate FEC, it should also possess FEC capabilities in the first place. Therefore for an accurate representation, the presence of FEC should be demanded as a contract, *and* the ability to control it as a knob. Another example of a related knob and contract would be in a CR MAC requiring control of an Non-Contiguous Orthogonal Frequency Division Multiplexing (NC-OFDM) transceiver. It would require a contract with the PHY transceiver to be *capable* of NC-OFDM transmissions, and a knob to allow it to control them, presumably by means of inputting a subcarrier mask.

Note that every CR MAC has potentially different PHY requirements, and in some cases multiple PHY transceivers may be needed. The Anatomy Schema acknowledges this, and each additional transceiver may be indicated on the Schema as necessary. Note that each transceiver possesses its own contracts, and its own knobs and meters, as are required by the Anatomy elements connected to it.

3.2.1.4 Awareness

Once a CR MAC protocol has been divided into its constituent elements, at least some of them should play a part in its awareness. Therefore, the Anatomy categor-

---

[3] Knobs and meters are terms commonly employed by researchers from Virginia tech to describe a CR's flexibility. Another example use of them is by Rondeau [6].

ises elements under "awareness" in addition to the four responsibilities re-examined in §3.2.1.1.

Awareness is in simplest terms information, gathered in order to inform decisions made by the CR MAC. It may comprise *external awareness*, that is to say information relating to matters outside of the CR device and the group of CRs of which it is a part. And as importantly, it may comprise *in-system awareness* relating to information about the CR device itself, and its surrounding group of compatible CR devices.

These two kinds of awareness are the two rightmost columns on the Anatomy Schema, and elements may be categorised under them both in the context of the spectrum resource allocation and exploitation roles of a protocol.

Considering the complexity of gathering and making use of sufficient awareness information in a CR, it is perhaps of no surprise that this particular aspect of the technology could be regarded as nascent. But it is important to recall that as discussed earlier, this is awareness of all resources available to a CR MAC.

3.2.1.5 Connection to higher layers

A completed Anatomy Schema shows the paths of information-flow between CR MAC elements. This information includes communications data passed down from higher layers to be transmitted over the air, and data received from the PHY to be sent up to higher layers following processing. Therefore, connections to and from the higher layers must be indicated as part of an Anatomy Schema. To fit within space constraints, these connections are placed in the bottom-right hand corner of the Schema, as in Fig. 3.1.

There are three journeys that will result in a complete Anatomy Schema. The first is the creation of one corresponding to an existing protocol. The second is one that results from the design of a new protocol from scratch, either comprising entirely new Anatomy elements, recomposing existing elements or a combination of existing and new elements. The third is one that results from the enhancement of an existing protocol by the addition of extra Anatomy elements. The remainder of the section will discuss each of these journeys in turn.

3.2.2 Creating the Anatomy Schema for an existing CR MAC

The steps involved in creating an Anatomy Schema for an existing CR MAC protocol are summarised in the flowchart in Fig. 3.2. The steps are represented pictorially in Fig. 3.3. The remainder of this section will discuss them in detail.

3.2.2.1 Step 1: Determine the role(s) played

The first step in creating an Anatomy Schema for an existing CR MAC protocol is to determine whether it plays a spectrum resource allocation role, a resource exploitation role or both. To do so it, is necessary to examine the protocol's definition of a "node", the network unit that it considers to be its scope.

A resource allocation protocol defines a node as an AP, a transmitter-receiver link, or a cluster of devices. It will not consider operations in terms of transmitter and receiver, and does not indicate how the spectrum resource should be employed within the node once allocated to it. In a case where a proposed protocol is unclear about its definition of a node, a lack of discussion of a channel access method – the precise decision of when to transmit – is an indication that the CR MAC protocol focuses on playing a resource allocation role.

A CR MAC protocol playing a resource exploitation role considers devices individually, transmitters and receivers, and should fully define the mechanisms of communication between them. Any CR MAC protocol that details both transmit and receive-side operation, detailing how a message is transmitted from one device and received safely at another, plays a spectrum resource exploitation role.

3.3 Completing an Anatomy Schema

3.3.0.2 Step 2: Determine PHY transceiver count

The next step is to determine the number of transceivers required by the protocol. This will usually be clearly stated at the outset in a CR MAC protocol proposal, but may not be, as the need for a dedicated transceiver for a Common Control Channel (CCC) is sometimes brought up only in the context of the discussion of the CCC.

Each individual transceiver will be marked on the Anatomy Schema.

3.3.0.3 Step 3: Indicate information path(s) and monolithic protocol form

This step involves considering the protocol in a monolithic form – as a single large element – encompassing all the responsibilities and awareness columns of whatever role(s) it plays. Then, information paths between the higher layers and each PHY transceiver must be indicated. Every transceiver will need some sort of connection, direct or indirect, to higher layers. The precise nature of this connection will be determined later in the process.

3.3.0.4 Step 4: Examine responsibility and awareness handling

This next step is among the most substantial. Each responsibility on the Anatomy Schema must be considered, and the CR MAC protocol examined in terms of if and how it meets it.

The answer to the question of how each responsibility is met – if it is met – is a

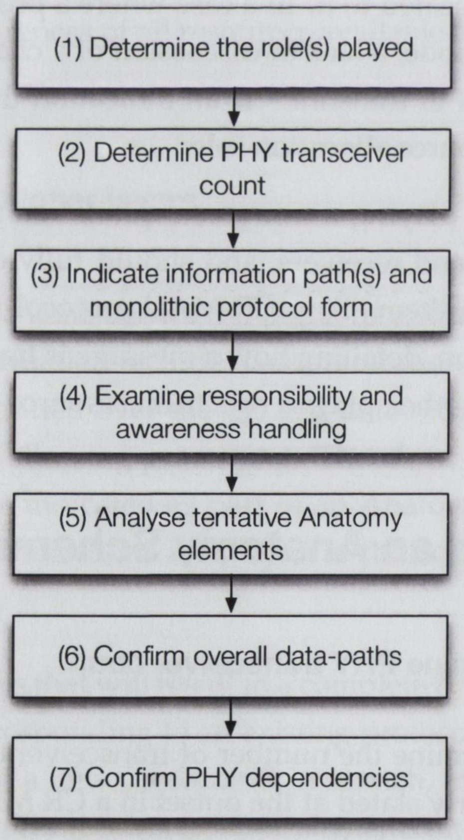

Figure 3.2: Flowchart outlining the steps involved in creating a complete Anatomy Schema for an existing CR MAC protocol.

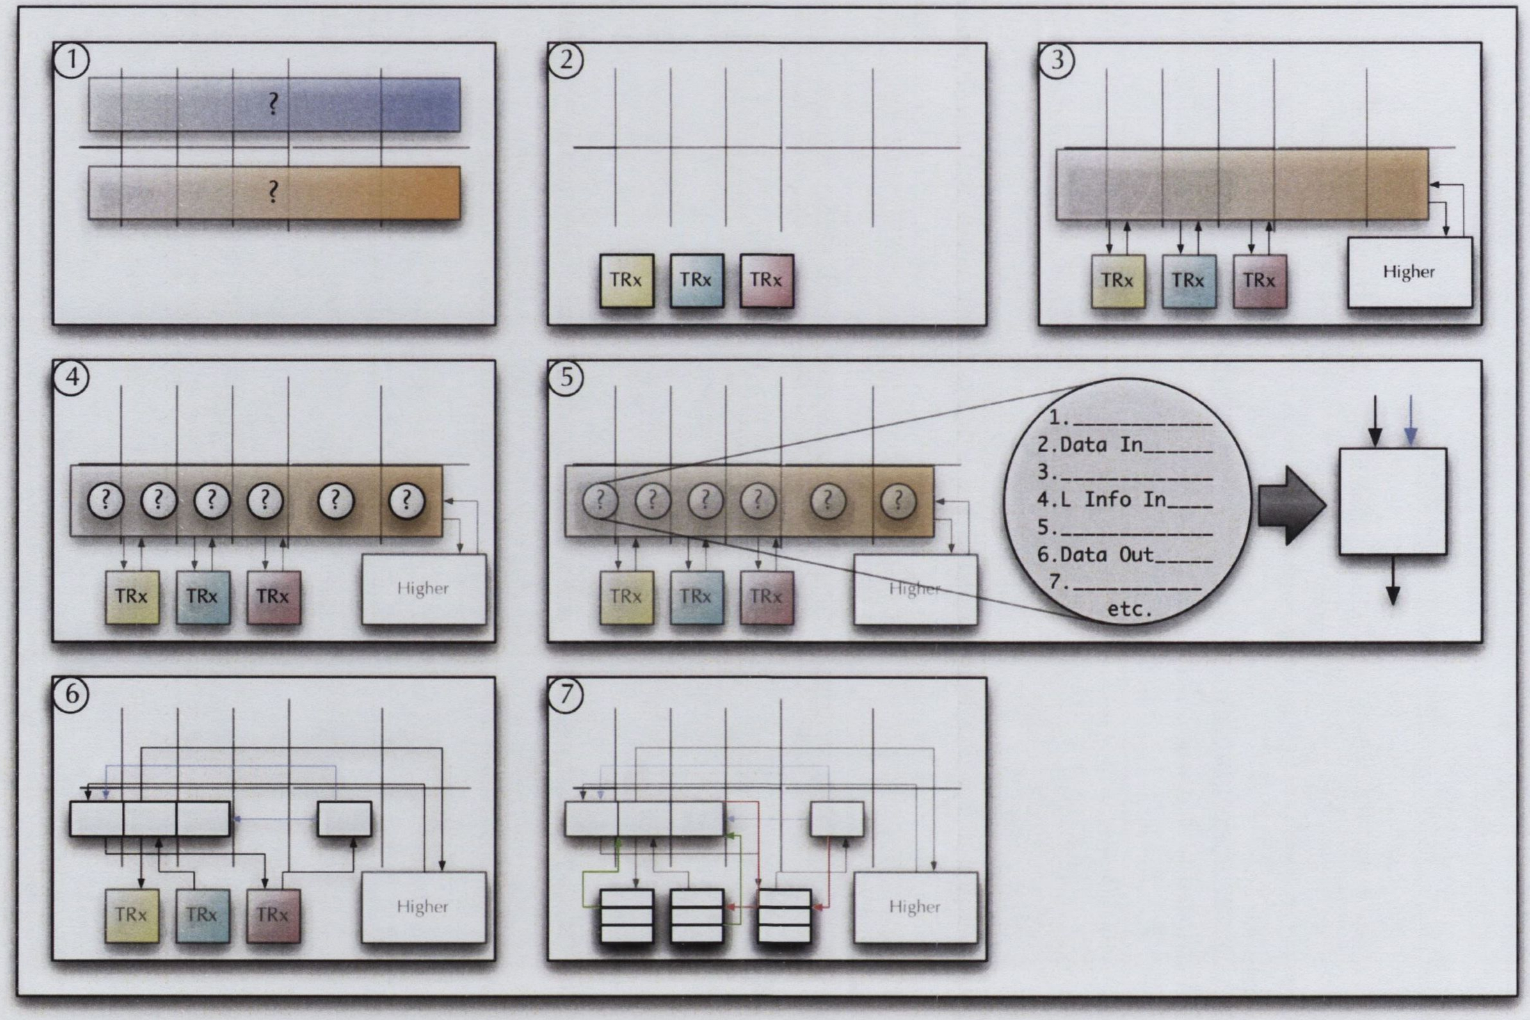

Figure 3.3: Pictorial representation of the steps shown in Fig. 3.2 to create the Anatomy Schema for an existing CR MAC protocol.

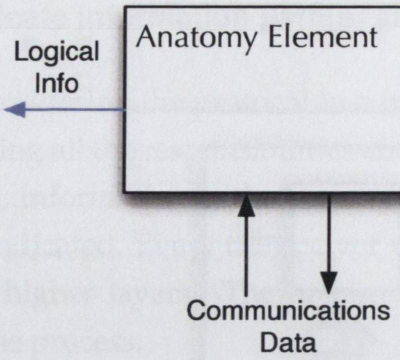

Figure 3.4: Simple representation of an Anatomy element, with data input and output paths, input of some logical information and output of some other logical information.

tentative Anatomy element. If there are potentially multiple answers to the question, there may be multiple tentative elements for that responsibility. This possibility will be returned to and discussed later, in §3.4.4. Additionally, in order to make the process clearer, and this step in particular, the creation of an example Anatomy Schema for a chosen CR MAC protocol will be explained in §3.3.0.8.

3.3.0.5 Step 5: Analyse tentative Anatomy elements

Now each tentative Anatomy element must be analysed, to confirm it is self-contained, and establish its inputs and output information paths. Anatomy elements can require input alone, or require input and output. Under the umbrella of information paths, there are two main categories. The first is *communications data*, which is any stream of bits that is destined, possibly indirectly, to be transmitted by the PHY. Conversely, it may also be any bitstream that was received from the PHY. In short, it is any path that eventually starts or ends with a transceiver. The second main category of information path is for when elements rely on information from other elements, not directly related to the communications data-stream. This might be the result of processing performed by an element, the occurrence of some event at an element, or information obtained by some other means relating to internal or external awareness. This form of information is termed *logical information* in the Anatomy, and encompasses any inter-element link not having to do with the communications data-path. An example generic Anatomy element is shown in Fig. 3.4, with communications data-paths indicated in black, and logical information paths indicated in blue. It should also be noted that there is a further special form of logical information input and output to an Anatomy element – its PHY dependencies, should it possess any. These are also analysed as part of this process, and then confirmed as part of a later step.

To analyse each element, it must be considered in detail, as a state machine or series of sequential steps. As long as the steps involved in the tentative element make up a closed loop and can completely define the addressing of this responsibility or part of the responsibility, then it can be deemed an Anatomy element. At each step, any interaction with or requirement for communications data input, output or for logical information should be noted. This list will make up the input and output of the element.

As an example, the Anatomy element corresponding to an ARQ mechanism will be considered here[4]. Teased out, the steps involved in the ARQ process are as follows:

```
1  When packet to send                   ▷ Data in from higher layers
2  Store copy
3  Send packet                           ▷ Data out to lower
4  while timeout == false
5       do if ACK received with same id  ▷ Data in from lower
6             then goto 1
7  if timeout
8     then goto 3
```

By observation, these steps appear to be a complete representation of the ARQ, and their only outside interaction consists of the the data output below at line 3, and data inputs from above and below at lines 1 and 5 respectively. Therefore the element to represent the ARQ need only have two communications data inputs and a single data output, and has no logical information requirement. Note that any PHY dependency/ies in an element will also be discovered as part of this process, and should be noted to be considered properly later, in step 7 (See §3.3.0.7.)

In the case of ARQ or other well known parts of a MAC or CR MAC protocol, it is not necessary to go through this process in detail every time a protocol is being deconstructed. Once the element characteristics are known, they can be reused.

---

[4] Recall that ARQ was explained in Chapter 2, §2.2.1.2.

3.3.0.6 Step 6: Confirm overall communications data-paths

With all elements confirmed in terms of their inputs, outputs, and logical information requirements, the next and penultimate step is to route the data-paths through all necessary elements, and draw links to represent any logical information requirements at components. In order to decide what order data-paths pass through the CR MAC protocol, the progression of the protocol must be followed, following the path a packet takes from the higher layers through the protocol. It should also be considered that there will commonly be multiple communications data-paths, especially if there are multiple transceivers. It is possible that some of these communications data-paths carry data that does not originate from higher layers, and is not bound for the higher layers of other devices. Control packets, for example are sent only from an element in one instance of a MAC, to terminate in the element of another. It is useful to distinguish further between these two forms of communications data. *General data communication* will be the term used to describe the communication of Service Data Units (SDUs) passed from higher-layers to be sent to corresponding higher-layers in other devices. **Control data communication** will indicate communications data generated solely by the MAC to facilitate execution of its roles and responsibilities, i.e. communications data internal to the MAC protocol, that will never be passed to higher layers. For clarity, Fig. 3.5 lays out the relationship between all categories of information-paths.

Note that while step 3 (§3.3.0.3) marks an information path between the higher layers and every PHY transceiver, this path need not be a communications data-path – it could be logical information. In fact, in the case of a protocol playing solely a resource allocation role, there will rarely be a communications data-path from the

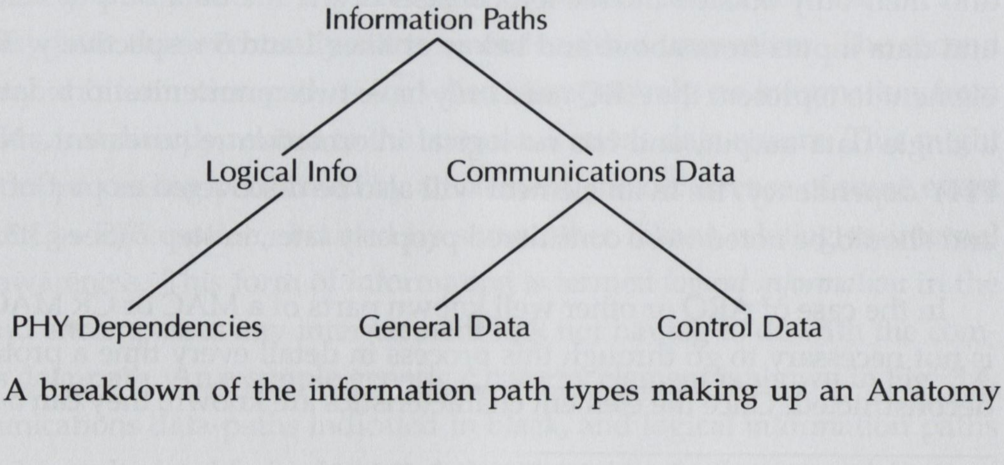

Figure 3.5: A breakdown of the information path types making up an Anatomy Schema.

higher layers to the PHY through the MAC's elements. This is because the spectrum allocation role involves only assigning spectrum resources to nodes – be they APs, links or device-groups – not detailing how they use it. It may control parameters of the PHY transceiver(s), but is unlikely to directly interact with the data being passed to it. A CR device whose MAC only plays a resource allocation role could therefore be considered incomplete, as exactly what happens to the SDUs coming from higher layers is not clear. In a case where a protocol playing only an allocation role is to be represented on the Anatomy for implementation, it is therefore useful to place a single element to play the simplest exploitation role. This need do no more than simply pass any general data it receives from higher layers to a PHY transceiver for immediate transmission on whatever resource is currently allocated to the device. Fig. 3.6 shows what such a simple element might look like on the Anatomy Schema of an otherwise complex allocation protocol.

The actual logic of such an element may be as simple as the following three steps:

```
When packet to send                ▷ Data in from higher layers
Attach address if needed
Send packet                        ▷ Data out to lower
```

However, should some-degree of reliability and performance be desired, then a more feature-rich exploitation implementation must be chosen.

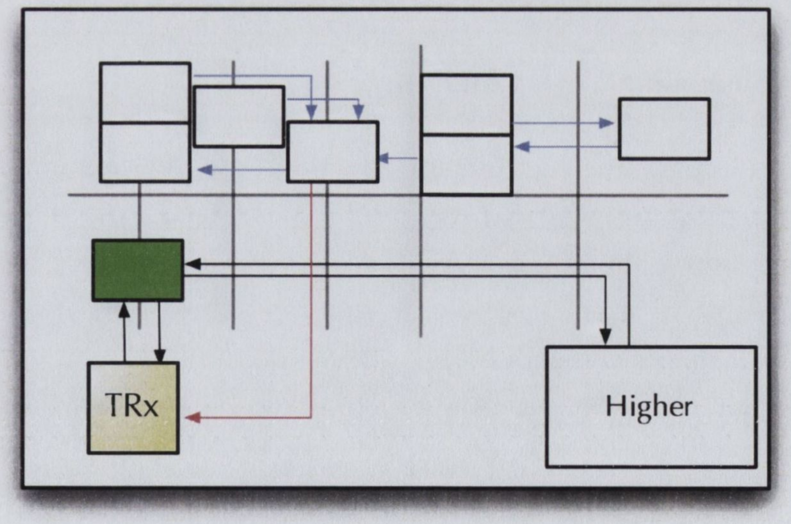

Figure 3.6: Overview of an Anatomy Schema for a protocol playing only a spectrum resource allocation role. It depicts a very simple exploitation element, highlighted in green, to make the communications data-path clear.

Any information path between the PHY transceiver(s) and the allocation elements is likely to be solely a PHY dependency to action or inform the allocation decision. Or it may be a control data-path somehow necessary to the resource allocation mechanism, such as for a CCC the allocation role requires.

3.3.0.7 Step 7: Confirm PHY Dependencies

Finally, it's important to consider first of all whether any elements have an implicit contract with a PHY transceiver or transceivers, demanding certain behaviour or capabilities of them. In order to deduce the requirement for a PHY contract, a transceiver must be considered to be a blank slate, of which no assumptions may be made. Data will be passed to it, to be transmitted as the transceiver wishes. Any requirements as to a specific mode of transmission, speed of processing, information validity and so on, are to be considered as out of the ordinary and marked as a PHY contract.

Any knobs and meters required of a PHY transceivers will have been noted as inputs and outputs of elements as part of the analysis process in §3.3.0.5. These dependencies must be noted on the relevant PHY transceivers, and linked to the dependent Anatomy elements on the Schema. See Fig. 3.7 for an example Anatomy element shown connected to the PHY knobs and meters on which it depends.

All of these steps were summarised in pictorial form in Fig. 3.3.

3.3.0.8 Example: Creating an Anatomy Schema

This section will apply the method introduced above, to create a complete Anatomy Schema representing a CR MAC protocol selected from the literature. The protocol in question is the "DCR-MAC" proposed by Yoo *et al.* [36]. This is chosen as it is a well featured protocol, drawn from a reputable publication. Furthermore, it has a variety of PHY dependencies, a need for multiple transceivers and multiple inter-component logical information requirements that allow demonstration of these parts of the anatomy. It is an opportunistic DSA CR MAC. The remainder of this subsection will give a brief overview of the protocol, before moving on to apply the steps outlined earlier and create an Anatomy Schema representing it.

DCR-MAC is a protocol for sharing the spectrum of multiple primary users. It assumes perfect sensing, which is considered by the authors to be the purvey of the PHY, and CR devices are controlled and co-ordinated using a dedicated CCC. They

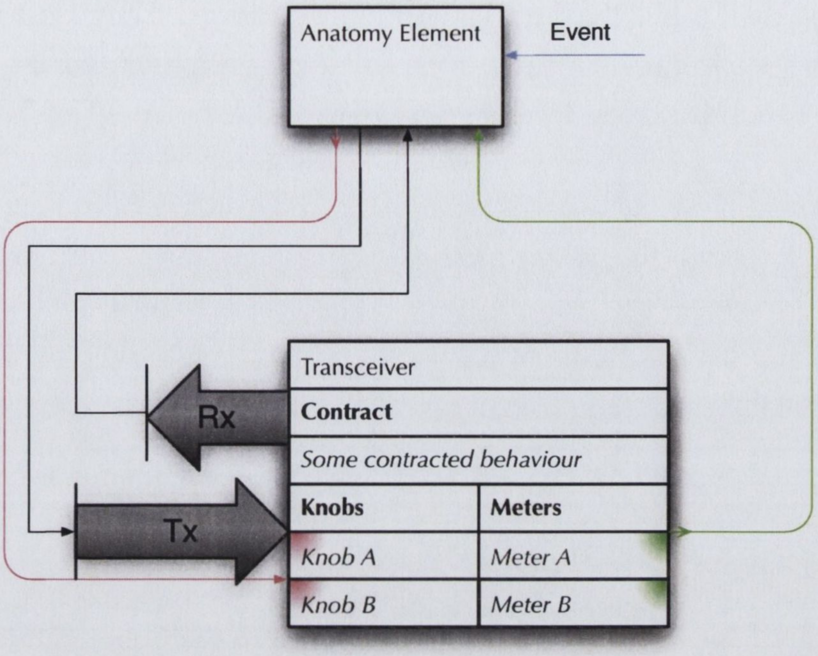

Figure 3.7: An Anatomy element shown with a PHY transceiver. The element has two PHY dependencies, on Knob B and Meter A, indicated by the red and green arrows respectively.

decide when and where to communicate using an extended version of CSMA/CA. A device wishing to transmit indicates so on the control channel using a Ready To Send (RTS) packet, which also contains a list of the channels known to be available to that device. In response, all neighbouring devices who are not the intended target of communication compare the channel list with their own. There follows a series of short time-slots, each one corresponding to a channel in the list provided in the transmitting device's RTS. If any neighbouring device knows of a primary user transmitting in one of these channels, it sends an unmodulated pulse in the time-slot corresponding to that channel. The transmitting device notes any pulses received, and checks if these indicate a change to the channels it had thought available. If there is any change, it issues an updated "RTSu" packet which contains the updated list. The desired receiver then issues a Channel Request, and waits for *its* neighbours to report channel occupancy in a similar series of timeslots, using pulses in the same manner. It then issues a Clear To Send (CTS) with an appropriately up-to-date channel list. The transmitter then broadcasts a Channel Reservation to its neighbours, and transmits for up to an agreed period of time. ARQ is the delivery assurance mechanism.

In this manner, what the authors entitle the "hidden incumbent node"[5] problem is solved. In addition, the handshake will alleviate conventional single-channel hidden and exposed nodes, to the extent that conventional RTS-CTS handshakes do. In this way, edge-cases are mitigated. There remains however the risk of the control channel becoming a bottleneck. Also, the expense of having the multiple transceivers necessary for such a protocol may be an issue[6]. Furthermore, though not explicitly stated by the authors, there is an implicit requirement for the devices to be synchronised in order to allow the short channel-reporting pulse time-slots to function as intended.

An Anatomy Schema corresponding to this protocol will now be created.

Step 1: Determine the role(s) played

The proposal for DCR-MAC clearly discusses events at both transmitting and receiving devices, fully defining the protocol required to arrange and carry out a packet exchange transaction. As it does so, it is clearly playing a spectrum resource exploitation role.

There is no discussion of the allocation of spectrum resources – channels or otherwise – in the proposal, therefore the protocol does not play a resource allocation role.

Step 2: Determine PHY transceiver count

The proposal discusses two possibilities: the option to employ two transceivers, one for "cognitive" data communication which communicates on idle frequencies allocated to PUs, the other tuned to a CCC to arrange this cognitive communication; or the option to employ a single transceiver, but use a split-phase control mechanism to arrange communication.

The two possibilities are in effect two discrete protocol choices. For this example, the former is chosen, as it is the default protocol explained in the paper. Therefore, two PHY transceivers are required. For the remainder of the process, the "cognitive" transceiver will be referred to as the device's *flexible* transceiver, the other its CCC transceiver.

---

[5] Identical to the standard hidden node problem as defined in §2.2.1, except that the hidden node in this case is a primary user.

[6] That said, it should be pointed out that a great many modern communications devices incorporate several discrete transceivers, and therefore this is not an unreasonable requirement.

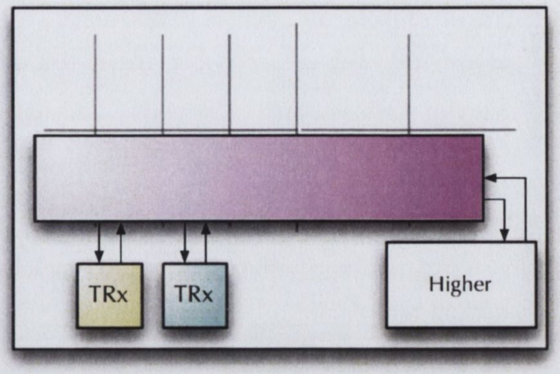

Figure 3.8: Sketch of "DCR-MAC" in simple monolithic form, indicating information paths from higher layers to transceivers.

Step 3: Indicate information path(s) and monolithic protocol form

With two transceivers, the simplest monolithic form of the protocol is as sketched in Fig. 3.8. Note the tentative information path to each transceiver has also been indicated.

Step 4: Examine responsibility and awareness handling

Traversing the Anatomy Schema left-to-right, the next step is to examine how the protocol handles responsibility and awareness. The column corresponding to "defining", what, when, and where the MAC accesses will require a tentative element relating to the "DCR-MAC's" transmission decision making mechanism. This is an extended carrier-sense mechanism, involving knowledge of other Secondary User (SU) activity, local PU activity and a multi-stage channel-confirmation handshake process to ensure the non-presence of PU.

In defining *what* the MAC accesses, the authors assume frequency has been channelised. Therefore the process of channelisation of frequency may also need to fall under this column.

The "deliver" column requires no more than an ARQ, as that is the only delivery assurance mechanism employed by the protocol.

The "control and coordinate" column requires a control-channel, as channel negotiation will occur over it.

The "edge-cases" column does not appear to have any specific tentative element. A major edge-case of concern in this particular protocol is the situation in which a receiving PU is a hidden node to an SU about to transmit. The end stages of the extended carrier-sense mechanism which has been located in the define column are already designed to minimise the risk of this occurrence however, therefore there is no specific entry for the edge-case column.

As regards its "external awareness", the DCR-MAC requires local information as to which channels are currently occupied by PUs, and which are vacant. The exact details of the mechanism used to obtain this information are considered out of the scope of the paper proposing the protocol. This non-specific information source is another tentative Anatomy element.

Finally, in terms of "internal awareness", the protocol employs virtual-carrier sensing to maintain knowledge of the activity of other SUs.

Step 5: Analyse tentative Anatomy elements

The next step is to consider the tentative Anatomy elements deduced above.

Again traversing left-to-right, the first tentative element is the extended carrier-sense mechanism. It will be termed "multi-stage carrier selection handshake" in order to fully describe it. It is the most complex element of the protocol, and therefore will require a detailed consideration in order to explain the element analysis process.The element must be broken down into the actions it performs step-by-step. Any step that requires interaction beyond the scope of the element itself, must be noted. As this handshake process varies at the transmitter and receiver side, it is broken down below, in both parts. The course of the handshake is now detailed. At the transmitter side:

```
When packet to send                        ▷ Data in from higher layers
Get local channel list                     ▷ Logical information in from
                                             PU-info and virtual carrier-sense
Send RTS with available channel list       ▷ Data out to CCC
Listen to neighbour PU status pulses       ▷ Data in from CCC
if neighbours update channel list
then
        Send RTS update                    ▷ Data out to CCC
When receive CTS                           ▷ Data in from CCC
Set PHY to indicated channel               ▷ PHY knob
Transmit Data                              ▷ Data out towards transceiver
```

And at the corresponding receiver side:

```
When receive RTS                             ▷ Data in from CCC
AND RTS and local channel lists              ▷ Logical information in from
                                             PU-info and virtual carrier-sense
Send Channel Request list                    ▷ Data out to CCC
Listen to neighbour PU status pulses         ▷ Data in from CCC
Update list with new status info
Send list in CTS                             ▷ Data out to CCC
Set PHY to indicated channel                 ▷ Logical information or knob to
                                             change channel
```

Summarising inputs and outputs, this deconstruction shows that the handshake element requires:

- Data in from higher layers
- Data in from CCC
- Data out to flexible transceiver
- Data out to CCC
- Logical information in from PU info source
- Logical information in from virtual carrier-sense
- Logical information out or knob to change channel

Therefore, the Anatomy element corresponding to it is as shown in Fig. 3.9. As it is now a confirmed element, it will be referred to in **bold** typeface for clarity.

The need for channellisation is the next item to analyse. The outputs of the **multi-stage channel selection handshake** element discussed above is the channel to transmit on, but a PHY transceiver does not usually operate in terms of channel, but rather of centre-frequency. Therefore, a **Channel mapper** element must maintain a look-up table of the frequency corresponding to each channel. It will receive channel choice as a logical input, and require access to a PHY knob controlling frequency, in order to select the appropriate one for transmission.

Examining the operation of the handshake element detailed earlier, a counter-intuitive observation can be made. Considering two handshake elements in two separate devices, arranging communication. Once the channel has been selected and set at the flexible transceivers, the receiving handshake element need not directly receive the packet itself. This can be handled directly by the **ARQ** element. With the inputs and outputs of this Anatomy element now clearly defined it is no longer tentative. The same process must be applied to remaining tentative elements in order to confirm their inputs and outs.

Continuing the process, the inputs, outputs and behaviour of ARQ were analysed earlier in the chapter, in §3.3.0.5, therefore they are known.

A **Control channel** element does little more than pass data through to a connected transceiver, therefore it also does not require detailed analysis. It will have a connection to the CCC transceiver, and any elements which make use of the CCC must have a communications data-path to it.

To represent the undefined source of information on PU spectrum occupancy, an element titled **Spectrum status** will be included. It is not well defined by the protocol, so its inputs are unknown and will not be marked. Its output will be logical information, to any elements requiring it. The **Multi-stage channel selection handshake** element discussed earlier requires such information, therefore it will require a logical information input from this element.

Finally, the **Virtual carrier-sensing** element will require data input from the control-channel, in order to listen in on channel selection handshakes. By means of these and the Network Allocation Vector (NAV) incorporated in the packets, it can build a map of which channels are instantaneously occupied, to provide to elements requiring it. Again, it is the **Multi-stage channel selection handshake** element which requires this information, therefore as mentioned before, it will also require a logical information input from this element.

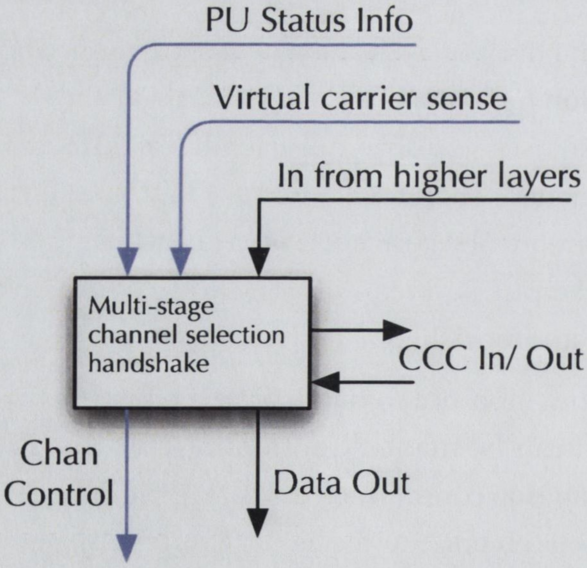

Figure 3.9: Anatamy element containing the **Multi-stage channel selection handshake** of the DCR-MAC protocol [36], with inputs, outputs and dependencies clearly indicated.

Step 6: Confirm overall data-paths

Inputs and outputs for all elements have now been clearly defined. The **Multi-stage channel selection handshake**, **ARQ** and **Control channel** elements require data input and output. The **Virtual carrier-sense** requires only data input. As the **ARQ**'s Data Acknowledgment (ACK) handshake occurs on the channel currently occupied by the flexible transceiver, it requires a direct input and output to it. The decision on when to transmit is made at the **Multi-stage channel selection handshake** element, therefore its data output is to the **ARQ**.

The CCC transceiver must be directly connected to the **Control channel** element, with both input and output. The **Multi-stage channel selection handshake** element will require input and output connection to this element, and the **Virtual carrier-sense** module will require input, in order to observe NAV information from control packets to maintain an up-to-date awareness of SU activity.

As stated earlier, the input of the **Spectrum status** element is unknown, and therefore for this Anatomy element it will be left unmarked.

Logical information connections are as discussed in the previous step – both the **Spectrum status** element and the **Virtual carrier-sense** element provide information to the **Multi-stage channel selection handshake** element.

The communications data-path connections are indicated, along with the PHY dependencies discussed in the final step, on the completed Anatomy Schema in Fig. 3.10.

Step 7: Confirm PHY dependencies

The PHY dependencies of the protocol must now be considered.

In terms of required contracts, two items of note stand-out in the protocol's specification.

The first is the unusual handshake mechanism, in which short energy pulses are used to indicate PU channel occupancy by neighbours. This level of fine-grained control is not necessarily a feature of all PHY transceivers. While in the case of many narrowband transceivers, the transmission of a very small number of Bytes by a MAC protocol would produce a correspondingly short pulse as its output, a transceiver using a technique such as Orthogonal Frequency Division Multiplexing (OFDM) parallelises the data being transmitted, involves extra overhead such as OFDM-frame preambles, and may potentially have a minimum packet size. What this means is the relationship between size of the input provided by a MAC protocol

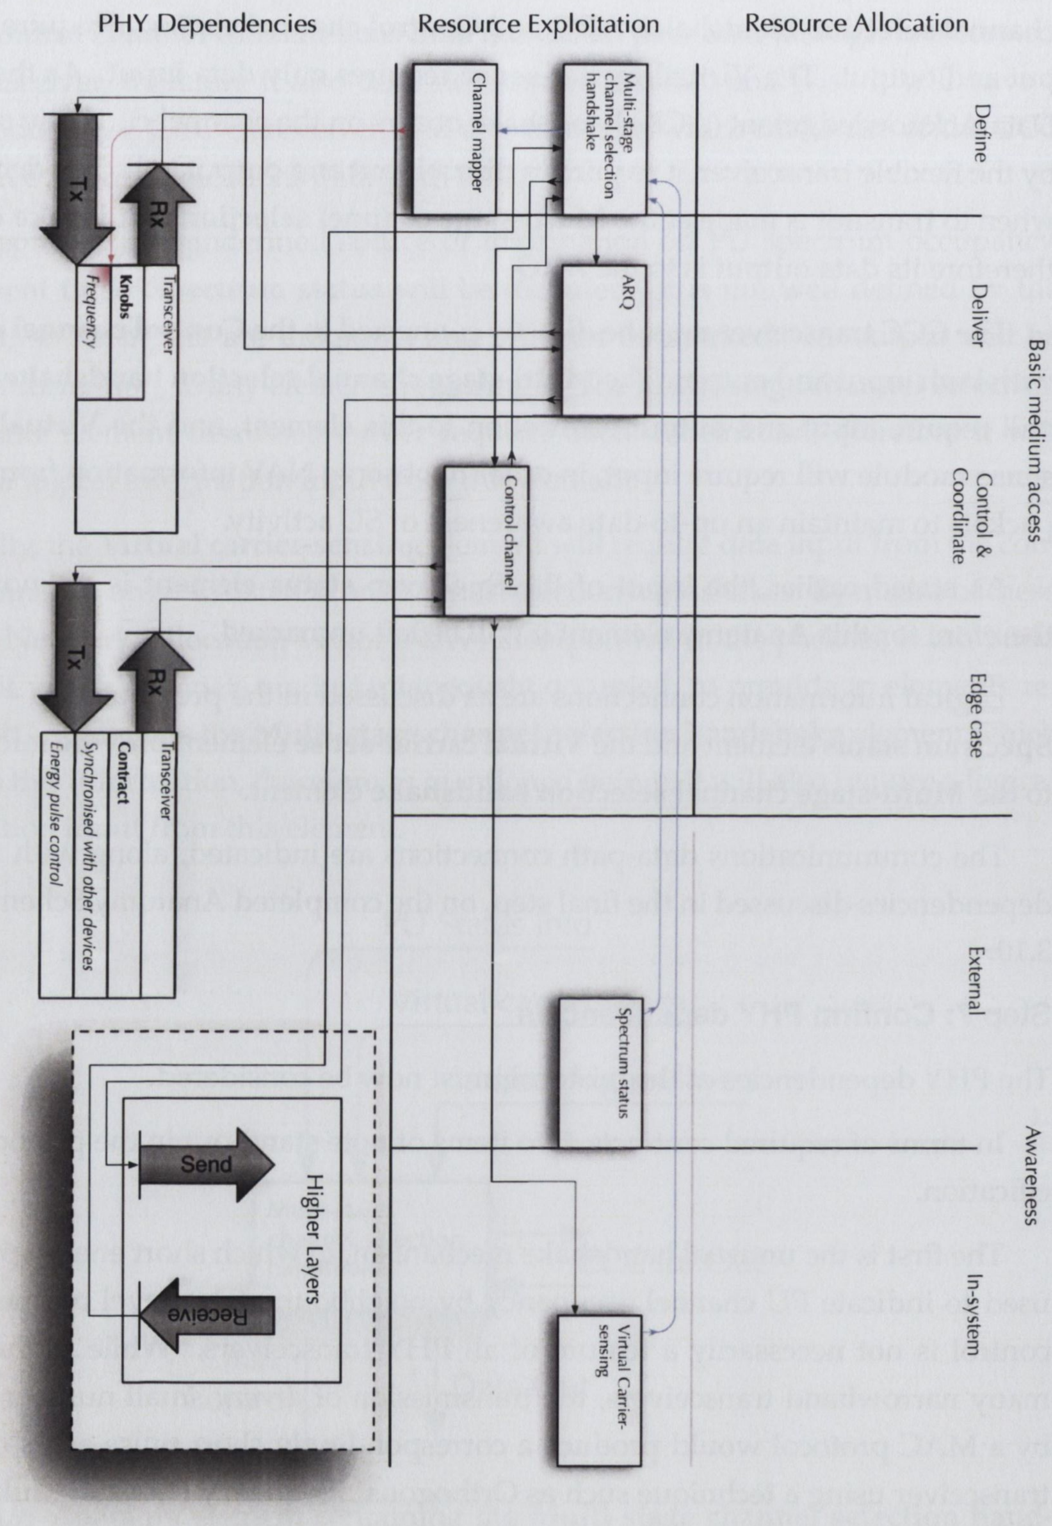

Figure 3.10: Completed Anatomy of the DCR-MAC proposed by Yoo *et al.* [36].

and the number of over-the-air symbols generated at the PHY is not necessarily one of direct proportionality. This therefore necessitates a contract.

The second item of note is that devices employing the DCR-MAC infer which frequency channel is being described by neighbours by the time at which energy pulses are received. They slot time, and map each timeslot to a frequency channel being described. For time-slotting such as this to have consistent meaning across devices, a time-synchronisation mechanism is required. As no mechanism is provided as part of the MAC protocol to arrange synchronisation, it must be expected of the PHY transceiver as a contract. Note that both of these contracts relate to the CCC transceiver.

Now any required knobs and meters must be considered. The need for them is revealed as part of the element analysis process in Step 5. Recall in this case that the **Multi-stage channel selection handshake** element required a PHY knob in order to be able to set the frequency channel once selected. No meters were required.

With these confirmed, they can be indicated on the final Anatomy Schema is as shown in Fig. 3.10, now complete.

3.3.0.9 Complete deconstruction: Safely extracting Anatomy elements

Complete deconstruction involves the consideration of Anatomy elements independent of the protocol in which they were proposed. Once an Anatomy element has been analysed as per Step 5 in the Schema construction flowchart in Fig. 3.2, its inputs and outputs are clearly defined, and it can be considered alone. When doing so, it is important to note not only the direction of data connections (i.e. input or output), but also their intended target. This target may be a a control communications mechanism such as a CCC, a data transceiver, or the higher layers.

Consider the **Multi-stage channel selection handshake** element of the DCR-MAC protocol deconstructed earlier, and shown in Fig. 3.9. Following analysis of it in the previous section, it could be considered to have been completely dissociated from the remainder of the protocol. Provided it received appropriate inputs and outputs, it could be incorporated into an alternative protocol and would function similarly.

Similarly, a logical information requirement by an element is agnostic to how the information requirement is met – any alternative element or group of elements providing an equivalent output can be substituted without affecting the functionality of the element.

3.3.1 Reconstruction: How to build a new CR MAC

Taking advantage of the disconnect between elements on a grand scale, is *reconstruction*. Once elements have been carefully broken down, and their communications data inputs and outputs, logical information inputs and PHY dependencies all carefully noted, groups of elements can be combined to create a new protocol from scratch.

Given a specification or scenario, the behaviour of each element can be combined to create a new protocol to create a protocol with the desired overall properties. If a sufficient range of Anatomy elements exists in order to fully satisfy the specification, the only difficulty in this process is ensuring that communications data-paths passing through all components are sensible. In order to explain this, elements with must be categorised as those which communications data pass through, which will be referred to as *passthrough* elements, and those which only require communications data as an input, which will be referred to as *absorbing* elements.

The connection of absorbing elements is straightforward. These are components which require communications data only as an input, presumably in order to perform analysis or processing of some variety, to provide a logical information or PHY knob control. The primary requirement in their case is to ensure that the component receives data from the appropriate source. An example is the **Virtual carrier-sensing** element in the DCR-MAC discussed earlier. In order to overhear control packets of neighbouring CR devices, its data input must be from the CCC, rather than the flexible transceiver as might initially be expected. This connection can be seen in the Anatomy Schema in Fig. 3.10. The exception to this simplicity is in the case where an absorbing element is situated on a data-flow subsequent to a passthrough element, which will be discussed below.

In the case of passthrough elements, there is both a communications data input and output. Data that has been received by the element is then sent on following processing. Here, the primary importance is to consider the effect of whether the input is *transformed* in some way before it is output. This can be done by gating it – withholding it from output for a time or altogether, or by directly altering it. When output, the effects of the transformed data on the behaviour of any subsequent elements on the communications data-path must be considered. For example, an absorbing element performing processing on received signals in order to produce an output, such as the **Virtual carrier-sensing** element mentioned above, will probably require access to all received data from a transceiver. Therefore it should not be placed on a communications data-path subsequent to a transforming passthrough component, unless the resulting behaviour from the absorbing element is desirable. In the example Anatomy Schema therefore, the **Control Channel** element is an example of a non-transforming passthrough component.

It is also important to note that every PHY dependency and contract requirement introduced by an element to be included in the overall protocol must be satisfied by the attached PHY. It is conceivable that the lack of availability of required PHY features may narrow the possible range of elements that can be applied to a given scenario.

The process of reconstruction will be discussed in greater detail and demonstrated in Chapter 4.

3.3.2 Enhancement: How to add features to an existing CR MAC

Enhancement is a special case of reconstruction, in which additional elements are added to an existing CR MAC protocol.

No specific rules in addition to those involved in reconstruction apply. It should be noted in some cases this may not be a purely additive process. The simple addition of an element may cause a conflict with the required inputs and outputs of existing elements, possibly necessitating a small number of removals of the original elements of a protocol, or modification of them. The only additional concern is the need to redouble efforts to avoid issues caused by the placement of elements on communications data-paths without consideration for the effects of transforming passthrough elements. With communications data-paths already laid down by the body of the existing protocol, the risk of misplacement of an absorbing element on a data-path becomes higher.

3.4 Observations on this view of CR MAC

This section discusses some of the ramifications of the view of CR MAC the Anatomy takes.

3.4.1 Representing centralised CR MACs

It should be noted that as an Anatomy Schema considers a CR MAC protocol deconstructed into its functional elements, what appears conceptually to be a single CR MAC protocol may – when viewed through the lens of the Anatomy – be revealed to be a set of two or more related protocols. In the case of a centralised approach, the AP and client MACs may be radically different from one another in every sense, from their elements and transceivers to their information pathways. Each will need to be represented on a unique Anatomy Schema. More generally, when setting out to follow the instructions for completion of an Anatomy Schema outlined earlier in this chapter, it is important in the case of any protocol involving heterogeneous CR devices to note that more than one Anatomy Schema may be required to capture this heterogeneity. Furthermore, the multiple Anatomy Schemata involved may be very different in appearance and complexity. Fig. 3.11 illustrates this with a simple low-resolution representation of the possible differences between the master/ base-station and slave/ client sides of a centralised MAC protocol.

3.4.2 Resource Allocation in real-world CRs

When viewed together on an Anatomy Schema, some issues relating to the separation of CR MAC protocols into the two roles of spectrum resource allocation and exploitation come to light. As alluded to in §3.1.2.3 earlier, protocols playing one

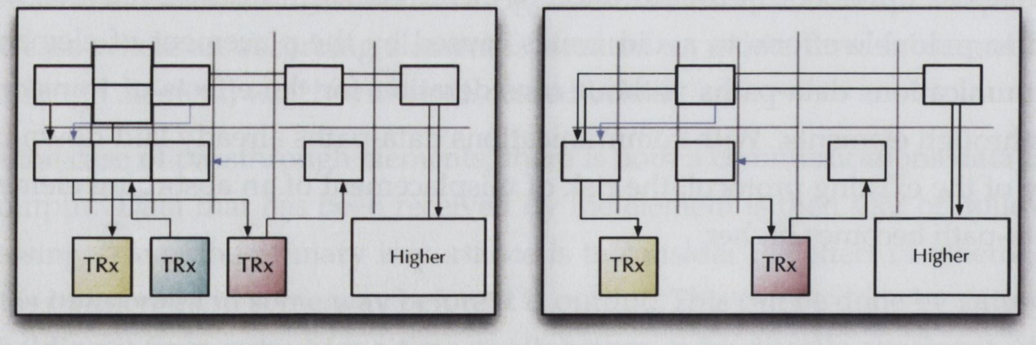

(a) Overview of a possible master/ AP protocol's Anatomy Schema.

(b) Overview of a possible slave/ client protocol's Anatomy Schema.

Figure 3.11: Simple representations of the Anatomy Schemata of the master and slave portion of a centralised MAC protocol. Note that the two may be very different in appearance, as may the work they perform and algorithms they deploy.

or other roles are almost exclusively proposed separately in the literature, and as a result, the precise difficulties relating to combining them are not entered into.

The problem relates specifically to the diverging view each protocol role takes of its scope. While exploitation deals with per-device, message-to-message control of access to a spectrum resource, allocation is concerned with the allocation of the resource, but not strictly with *what* the allocation is to. Resource allocation proposals speak of *nodes* being allocated to, but may variously define them as links, APs or groups of devices. In a real-world implementation therefore, for every group of devices, only one device is likely to be performing the allocation, while the remainder of the group abide by it. The question is, which device? This is somewhat more trivial in the case of a link or AP, where the decision of which device is allocating is intuitive. But for the group case, it is an issue unspecified in current protocols in the literature.

To fully address this issue, additional Anatomy element(s) should be responsible for i.) deciding the memberships of such groups[7]; ii.) deciding the allocating member, or "leader" of the groups and iii.) the distribution, communication or discovery of allocations by all members of the group. This is of course in addition to the requirement for elements to enact the allocation decision. How these three steps are performed is at the discretion of a CR MAC designer. However the fact that allocation and exploitation are generally dealt with separately in the literature means the two diverging world-views, one device-based, and one link or group-based, rarely meet. Given this lack of discussion of these combination issues in the literature, the remainder of this section will be devoted to discussing how appropriate solutions would be represented using the Anatomy.

These element(s) obviously form part of the Spectrum Resource Allocation role of the protocol's Anatomy Schema, and excepting in unusual circumstances fall under the responsibility to Control & Coordinate. The degree of functionality required depends on whether the protocol is centralised or distributed.

As noted in the previous section, in the case of a centralised protocol, an Anatomy Schema representation views the MAC as two separate protocols. Therefore in such a case, allocation decisions will usually take place at the AP or base-station. No specific elements are required to determine which is the group leader, but at the discretion of the protocol designer, there may be one or more involved in determining group membership, i.e. determining which clients belong to this AP. There will

---

[7] It should be noted that group-formation and cluster-formation are well-studied problems, and the literature contains many proposed solutions that could be deployed [82].

also be a need for an element or elements responsible for disseminating the allocation decisions to client devices. Considering the client-side Anatomy Schema, the only element(s) required for spectrum resource allocation are to determine group membership, i.e. which AP this client is attached to. This may be an extremely simple, or even fixed mechanism. There must also be some element responsible for receiving and taking any necessary actions in response to the allocation decisions made by the AP or base-station. This is in addition to whatever elements are involved in making and acting on the allocation itself. Figs. 3.12a and 3.12b provide a simple representation of how matching AP and client protocols' allocation elements might differ.

In the case of a distributed protocol complexity is increased. All three of the points mentioned earlier must be actively addressed by Anatomy element(s): group membership must be determined, group leadership must be determined, and the information both as to which device allocates and what its allocation decisions are, must be disseminated. For the protocol to be truly distributed, every device must be *capable* of allocating, and so each must possess all elements involved in the protocol's spectrum resource allocation role. In situations where the device is part of a group but not allocating, these elements will simply not be in use. Note that as they are all involved in the allocation mechanism, there is likely to be a significant degree of dependency between them. Alternatively, a combined element may be responsible for all three of these points. Fig. 3.12c depicts an example simplified Anatomy Schemata showing how a complete distributed CR MAC protocol's allocation would be represented.

The above examples are general enough to be representative of the combination of spectrum resource allocation and exploitation roles of many protocols. In particular, they highlight the extra elements that will be required in cases where a protocol proposed as only playing an allocation role is being incorporated into a system. The anatomy by no means precludes the representation of more complex CR MAC protocols, however. For example, it may be the case that some spectrum resource allocation mechanism may not have one one effective leader, but that groups make and disseminate collaborative allocation decisions.

The failure to address this issue is a limitation of the proposals made in the literature thus far, rather than being a limitation of the Anatomy of a CR MAC. Therefore for the remainder of this thesis, in cases where protocols playing only a spectrum resource allocation role are discussed, the following solution is proposed: group membership is deemed to have been externally chosen in advance and known by all devices. Group leadership is deemed to be granted by means of a simple election mechanism[8] and all remaining devices in a group are deemed to have disabled their spectrum allocation portions, save that part necessary to receive allocations from their local group-leader. This will not be depicted on the Anatomy Schemata, but could be done so in a manner similar to the examples discussed in this section.

---

[8] An election mechanism will be presented later, in Chapter 4, §4.2.4.4, for the alternative purpose of deciding a temporary master in a cluster-based Time Division Multiple Access (TDMA) system. This mechanism could however also be applied to select an allocating device in a group.

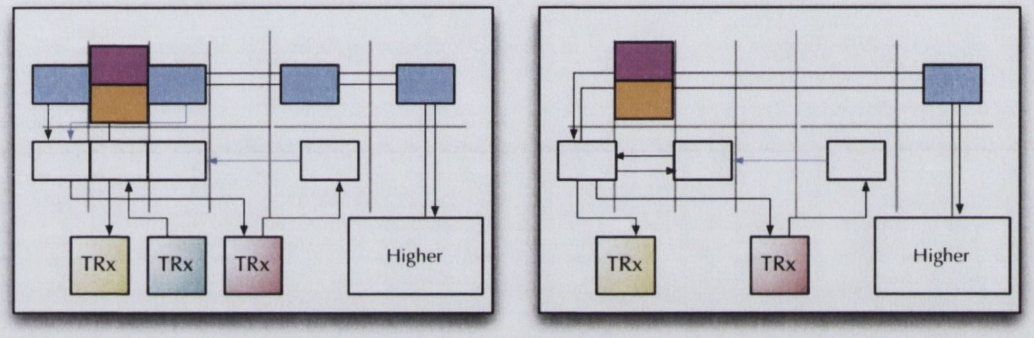

(a) Overview of a possible master/ AP protocol's Anatomy Schema, with elements required to support allocation in a real-world system highlighted.

(b) Overview of a possible slave/ client protocol's Anatomy Schema, with elements required to support allocation in a real-world system highlighted.

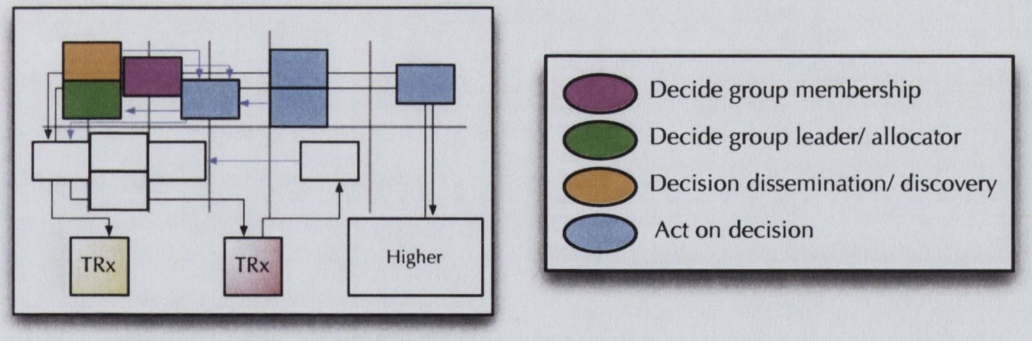

(c) Overview of a possible distributed protocol's Anatomy Schema, with elements required to support allocation in a real-world system highlighted.

(d) Key as to roles played by allocation elements in Figs. 3.12a-3.12c.

Figure 3.12: Simple representations of how the Anatomy Schemata of the master and slave portions of a centralised MAC protocol, and of a distributed protocol, might look. Elements necessary to support real-world allocation are highlighted. Note the relative increase in complexity for a distributed protocol.

3.4.3 MAC Dependencies

The importance of being aware of PHY contracts, knobs and meters is discussed above, and their inclusion is a significant contribution of the Anatomy. In the same way that CR demands a violation of the MAC-PHY boundaries to a significantly greater extent than even violations which assisted in the implementation of wireless communication on systems initially designed for fixed wireline traffic, it is more than reasonable to assume that such violations will be expected between other layers. It is misleading to suggest that cross-layer research is exclusively the purvey of the combination of MAC and PHY. Therefore in light of the significant research in the area of Cognitive Radio networking, considering the effect of cognitive radios on a wider scale, it is wise to assume that cognitive networking protocols will require the same sort of exposure of direct interfaces and the provision of contracts by their MAC as that MAC does of its PHY. Given that, it is easy to expose similar MAC dependencies for a set of network and higher layers. These would be marked as part of the higher layer connection on an Anatomy Schema, and require such Anatomy elements as required to provide their capabilities.

3.4.4 A note on subjectivity

A final item of note relates to constructing an Anatomy Schema. It should be mentioned that the process allows sufficient variation that some degree of subjectivity may appear to be involved in the process. It is however more accurate to say that the choice of granularity of elements means that there are a multitude of accurate Anatomy Schemata for a given CR MAC protocol, and the choice of granularity with which to represent it is at the discretion of its designer and/ or analyser. In particular, when performing comparisons between multiple protocols, it may be important that both be considered at the same relative granularity.

It should also be mentioned on the subject of subjectivity that the placement of elements on an Anatomy Schema under the columns corresponding to responsibilities and awareness is not a strict matter. By contrast information paths, encompassing element input/ output and PHY dependencies is very much so, and must be carefully indicated to characterise protocol behaviour correctly.

3.5 Summary

This chapter detailed a number of concerns related to CR MAC research, both as performed thus far, and going forward. In response to these concerns, it introduced the Anatomy of a CR MAC, a novel means of accurately representing a CR MAC protocol in terms of its constituent functional elements. It explained a general method for creating the Anatomy Schema corresponding to existing protocols and also how the Anatomy elements of these protocols could be considered alone, provided dependencies both on other elements and on the PHY are taken into account. It continued to discuss the manner in which Anatomy elements resulting from a deconstruction of existing protocol could be combined to create wholly new CR MAC protocols, and integrated into existing protocols in order to provide additional functionality. And it briefly covered the ways in which the Anatomy enhances CR MAC research, both existing and in the future.

The following chapter will demonstrate the Anatomy of a CR MAC "in action", by mapping its design method to a real world implementation.

3.4.3 MAC Dependencies

The importance of being aware of PHY contracts, know and [illegible] is described above, and their inclusion is a significant contribution of the Anatomy. In the same way that CR demands a violation of the MAC–PHY boundaries to a significantly greater extent than ever [illegible] in the conventional [illegible] communication on systems [illegible] designed for fixed wireless [illegible]. It is more than reasonable to assume that such violations will be expected between other layers. It is misleading to suggest that cross-layer research is exclusively the purview of the combination of MAC and PHY. Therefore in light of the significant research in the area of Cognitive Radio networking, considering the object of cognitive radio on a wider scale, it is wise to assume that cognitive [illegible] protocols will require the same sort of exposure of direct interfaces and the provision of contracts by their MAC as that MAC does of its PHY. Given that, it is easy to expose similar MAC dependencies [illegible] of the higher layer [illegible] Anatomy Schemas [illegible] elements as required to provide their capabilities.

3.4.4 A note on subjectivity

A final item of note relates to constructing an Anatomy Schema. It should be mentioned that the process allows sufficient variation that some degree of subjectivity may appear to be involved in the process. It is however more accurate to say that the choice of granularity of elements means that there are a multitude of accurate Anatomy Schemata for a given CR MAC protocol, and the choice of granularity with which to represent it is at the discretion of its designer and/or analyser. In particular, when performing comparisons between multiple protocols, it may be important that both be considered at the same relative granularity.

It should also be mentioned on the subject of subjectivity that the placement of elements on an Anatomy Schema under the columns corresponding to responsibilities and awareness is not a strict matter. By contrast information paths, encompassing element input/output and PHY dependencies is very much so, and must be carefully indicated to characterise protocol behaviour correctly.

4 The Anatomy in action

This chapter will demonstrate the use of the Anatomy of a Cognitive Radio (CR) MAC to design and construct CR Medium Access Control (MAC) protocols. It is here implemented on the Iris software radio platform [83] in tandem with the Ettus Research Universal Software Radio Peripheral (USRP) front-end[1] [84]. It will begin by discussing the platform chosen for implementation, and the method of doing so. It will then continue by providing an example implementation, using Iris, of a cognitive MAC, consisting of allocation and exploitation elements, and having multiple physical (PHY) dependencies. In doing so, it will take care to discuss how elements created in the implementation might be reused, discuss how the use of the Anatomy in the design model allows for easy future expansion of features.

4.1 A platform for Implementation

The Anatomy is a tool for the understanding and design of CR MAC protocols. In order to take advantage of the design philosophy it espouses however, a platform for implementation must be chosen. Several experimental platforms for CR research exist, including GNU Radio [85], the University of Rice Wireless Open-Access Research Platform (WARP) Project [86, 87], Papyrus [88] and many others. They differ in terms of their processing model, the level of abstraction they offer the user, the degree of flexibility, and the steepness of their learning curve. Furthermore, some platforms operate on custom processing hardware, while others operate on conventional General Purpose Processors (GPPs). They also vary in terms of whether they integrate the radio front-end hardware and processing elements onto a single platform, or whether they make use of separately attached front-end hardware. Lotze

---

[1] It should be noted that Iris is not limited to use of the USRP as its radio front-end. It is merely the chosen equipment in the context of this implementation as it is a popular and widely available platform.

offers a more detailed discussion of various Software Defined Radio (SDR) platforms and their tradeoffs for, which the interested reader is referred to [79].

Though nothing explicitly prevents it, fewer experimental platforms appear to have been employed for multi-device research and CR MAC research. This is likely as in many cases their initial design was as PHY layer experimental platforms, and making the move to MAC layer research can in some cases stress existing design assumptions, or necessitate changes in the way data is passed through the system. To consider the examples mentioned above, GNURadio is a platform that was designed as a PHY research tool, and still sees limited research output at the MAC layer. Papyrus and WARP on the other hand are newer platforms, and were designed with MAC research in mind, so they possess feature-sets designed specifically to support it.

Iris is a SDR platform developed at CTVR in Trinity College [79, 83, 89]. It is a highly flexible and reconfigurable radio system, running for the most part on GPP hardware. It divides radio Digital Signal Processing (DSP) tasks into blocks of a relatively high level of granularity, but do not sacrifice flexibility. Owing to the extensive existing library of components already implemented, and to the body of knowledge built at Trinity during its development, Iris was chosen as an example platform for implementation of the Anatomy. Notwithstanding this choice, it should be stressed that any SDR platform is a potential option for implementation of the Anatomy's design philosophy.

4.1.1 Introduction to Iris

The following paragraphs provide a brief introduction to Iris, its design methodology and its features. Iris has been in development for several years, and as such is work built up with the assistance of many researchers. Therefore though the following material is not new, it provides an important introduction to the reader, revealing relevant details which must be understood as part of process of implementing the Anatomy of a CR MAC.

Iris implements radio functionality in a block-wise manner, and graphs of these processing "blocks" can be built to create one or more radio chains. In Iris these blocks are referred to as *modules*. Iris itself is divided into the *core* which is what is executed at runtime and is rarely modified, and these modules. There are two types of Iris module, `Components` and `Controllers`, but the most commonly used is the `Component`. All Iris modules are implemented in C++, after which they become

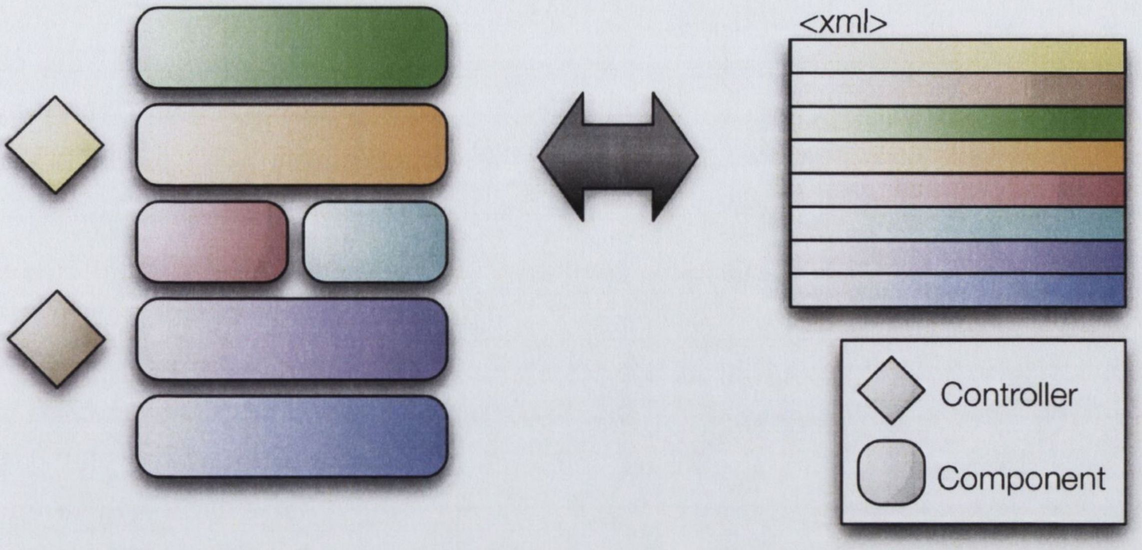

Figure 4.1: Graphical representation of the transformation of an Iris configuration of components and controllers into a corresponding xml file.

part of its library[2]. An xml configuration file is used to select the `Components` from the library to be used in a specific radio implementation, and to specify how they should be connected to create a radio chain or chains. A valid xml configuration is provided to the Iris core at runtime, which verifies it for consistency, instantiates the radio chain or chains as specified, and starts to run them. A representation of this transformation is shown in Fig. 4.1. Note that for the remainder of this thesis, figures representing Iris implementations will depict `Components` as round-edged rectangles, and `Controllers` as diamonds.

To facilitate reconfiguration of `Component` behaviour, `Component` designers expose *parameters* which can be set using the xml configuration, and may be changed at run-time. Communications data-flow, namely the streams of the bits or samples that make up the radio's *data-plane*, happens between `Components`, but there is another kind of information-flow in Iris. This second information type is *control-plane* information. It was mentioned above that Iris possesses two module types, and the second kind of module – termed a `Controller` – is what handles this control information, which is made up of `Events`, `ParametricReconfigurations` and `Commands`. Specifically, `Events` are signalled by `Components`. `Controllers` subscribe to and then respond to the activation of `Events` by performing their de-

---

[2] It should be noted that the term *library* here is used to mean the available array of implemented Iris modules. It does not refer to the compiled object code libraries which make up Iris, and are dynamically linked at runtime to be instantiated and run as a radio.

sired processing, or reconfiguring the parameters of `Components`. `Controllers` reconfigure `Components` by issuing `ParametricReconfigurations`. They can also issue `Commands`, which are control messages that `Components` can wait for as a signal that some event has occurred in another `Component`.

Iris is a platform which performs radio DSP, but it requires hardware in order to actually perform over-the-air radio frequency (RF) transmissions. Any hardware capable of accepting complex base-band signals as an input for transmissions is suitable. For implementation in this thesis, Ettus USRP hardware is chosen [84]. USRP hardware is a range of affordable SDR front-ends, which can be attached via standard USB or Ethernet to a PC, and accepts input in the form of complex base-band signals. Through the use of a selection of possible daughterboards, a wide spectrum of frequencies can be transmitted and received upon. Iris interfaces with the USRP by wrapping instantiations of its driver software in `UsrpTx` and `UsrpRx` `Components`, which act as sources or sinks in a radio chain. Parameters of these `Components` allow for initialisation and run-time reconfiguration of USRP parameters, such as frequency, rate/bandwidth, gain, and a number of features assisting transmission, detection, timing and scheduling of signals. It should be noted that the USRP is by no means exclusively compatible with Iris, and was initially conceived of as companion hardware for the GNURadio project. While a great asset to the research community, it should also be noted that the hardware suffers from a number of limitations, which will be discussed later in the chapter.

4.1.2 Iris: From links to networks

Previous Iris implementation and experimentation work focused on advanced PHY layer techniques. The experimental systems built were unidirectional continuous links, consisting of a single transmitter and receiver. With a move to communications between groups of devices, certain assumptions no longer hold. Multi-device communication requires *transceiver* operation, and is inherently packet-based, as no one individual device is likely to be transmitting the majority of the time.

The Iris library contains both narrow-band and wide-band modulator and demodulator `Components`. However, as a result of the link assumptions mentioned above, some of the existing `Components` were unsuited for MAC-layer development. The signal-processing `Components` used for synchronisation and timing-recovery of signals in order to properly demodulate them were simply not designed with discontinuous reception in mind. The existing Orthogonal Frequency Division Multiplexing (OFDM) `Components` were however capable of discontinuous reception, and so would only need some modifications to be useful for MAC layer research.

As OFDM is a multi-carrier wide-band transmission technique that is highly spectrally efficient and posseses great potential flexibility, it is already of significant interest in the context of CR. Therefore for the implementation discussed in this chapter, it represents an interesting and relevant choice. In their original form, Iris's OFDM `Components` assumed all data transmissions were to be of fixed and equal length. This restriction had to be removed. It was inefficient, given a MAC protocol is likely to send packets of many different lengths. Data packets are for example likely to be significantly longer than control packets. The OFDM components were modified to transmit PHY level header bits, which allowed variable OFDM frame length on a per-frame basis. With such a header in place, it was a logical further extension to include a Cyclic Redundancy Check (CRC) sequence in it. This would allow the demodulator to determine if a frame had been received correctly.

In a more general sense, as part of the process of modification, it was also necessary to experiment with the USRP's transceiver capability. This was done in order to determine the USRP's behaviour and reliability when combined with Iris in the case where multiple USRPs transmit and receive on the same channel. In some cases this caused the unexpected behaviour of some `Components`, which had to be corrected for. For example, a result of the rapid and constant switching between transmit and receive modes was that spurious data could be introduced into the stream of received samples passed by the USRP to Iris. Extreme values of this data, unlike any that would correspond to a real received signal, could intermittently cause some DSP operations such as those involved in OFDM preamble-detection to fail, and enter a permanent bad state. Such errors were corrected for in order to proceed with implementation.

Though these changes were made with a view to enabling implementation of the Anatomy of a CR MAC, such modifications would have been necessary to facilitate MAC layer research of any kind, in any system initially built primarily with link-only communications in mind.

4.1.3 Applying the Anatomy of a CR MAC to Iris

To implement a CR MAC using the Anatomy in Iris, the steps in Fig. 4.2 must be followed. This section will talk through the steps involved in this flowchart. It should

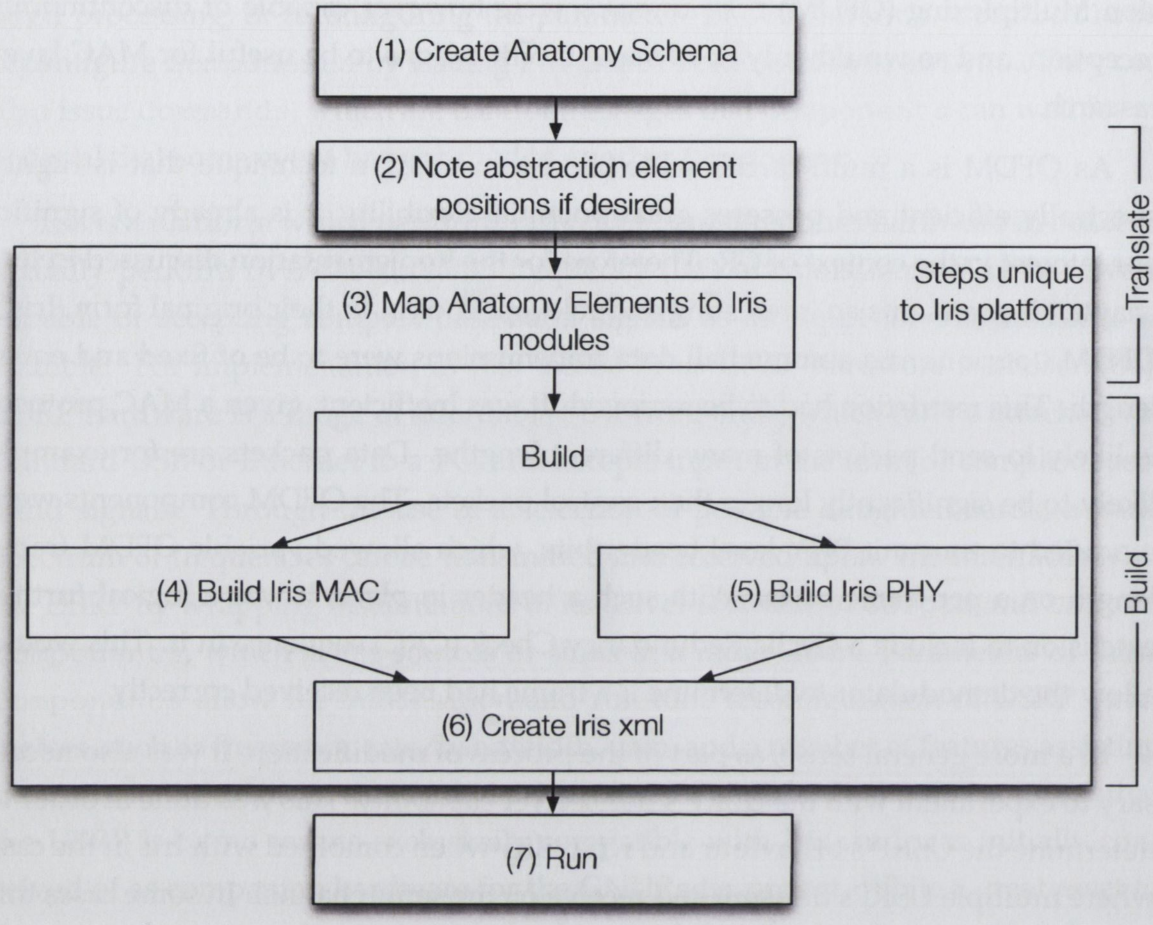

Figure 4.2: A flowchart detailing the steps necessary to implement a CR MAC protocol created in Iris, using the Anatomy.

be noted that the majority of the steps involved are specific to implementation in Iris. Nonetheless, implementation on another platform would not be significantly different, and would also follow the four overall steps highlighted in Fig. 4.2 above. These are: Create Anatomy Schema if it's not present already; Translate the model as required and desired for the platform; Build the protocol, implementing any functionality not currently available; Run and test it.

4.1.3.1 Create Anatomy Schema

The first step to implementing a CR MAC using the Anatomy is to create the Anatomy Schema corresponding to the protocol. To do this, the instructions detailed in Chapter 3, §3.3 must be followed.

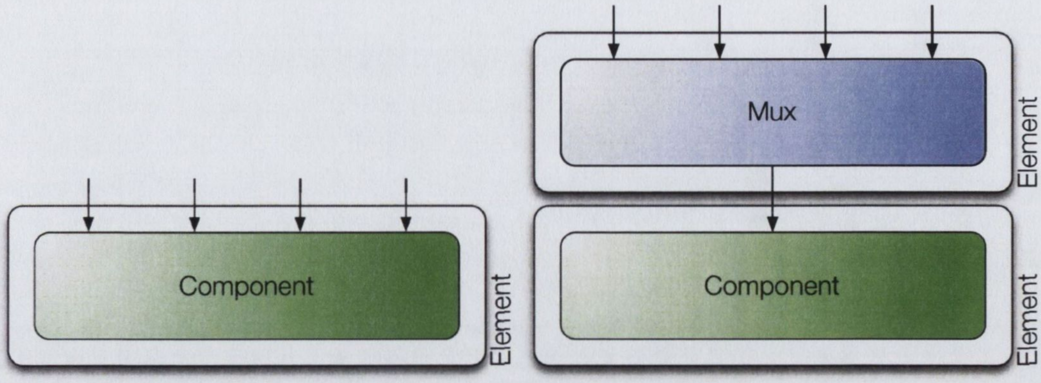

(a) Element implemented in Iris without use of an abstracting **Mux** element.

(b) Element implemented in Iris to use an abstracting **Mux** element. (Implemented in Iris as a `MuxComponent`.)

Figure 4.3: Examples of implementation in Iris without and with use of **Mux** elements. A change in the number of inputs in (a) will require modification to the `Component`. But in (b), a change in the number of inputs only requires replacement with an appropriately sized `MuxComponent`.

4.1.3.2 Note abstraction element positions

In order to facilitate evolvable implementation and allow data flow to be consistent, there are several points at which "abstract" elements may be useful. These can serve to multiplex or demultiplex communications data-paths, and to keep elements abstracted from direct dependency on one another. They are elements that are not represented by default on an Anatomy Schema, because they do not play a strict part in MAC protocol functionality, instead they serve as implementation aids. They tie elements together and act as consistent interfaces between them. In Iris they are implemented as `Components`, but as noted in the flowchart Fig. 4.2, they are not specific to Iris implementations.

Implementations in this thesis will make use of 3 abstraction elements. The first of these is the **Mux** element, which serves to multiplex and demultiplex communications data-paths. Use of **Mux** elements allows elements to be implemented independent of the future number of inputs. Fig 4.3 compares the implementation of an element in Iris, with and without use of a **Mux** element

A **data channel** element will be used to serve as an abstract representation of the endpoint of a general data-path for messages from higher layers to be transmitted. These are the messages containing the bits passed from the higher layers that must be transmitted over the air. Elements that output data to be sent towards a trans-

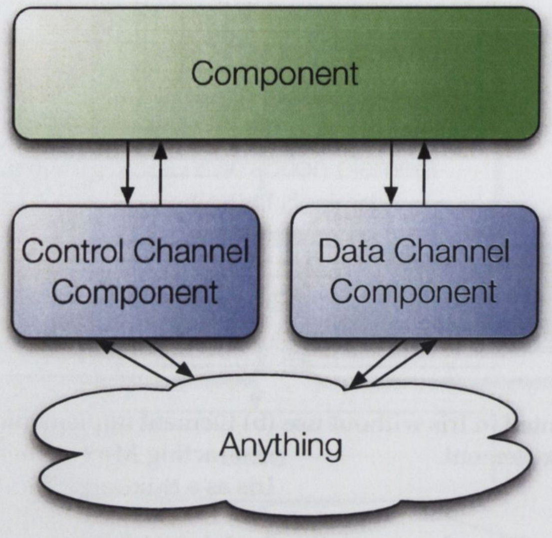

Figure 4.4: Representation of how two abstraction elements – implemented in Iris as `Components` – represent a Control Channel and Data Channel to separate an element from what it is specifically connected to.

ceiver can instead output data to this element, allowing the nature of the actual communication to be abstracted.

A **control channel** element serves a similar role, representing the endpoint of a control data-path[3] for control messages. A device that passes its messages to a **control channel** element need not be aware as to whether the element is connected to a dedicated transceiver, a scheduler sharing a channel with other communications data-paths, or handling control data messages by some other means.

These two elements serve to abstract other Anatomy elements from the actual implementation of the main data channel and the control channel. Used appropriately, they facilitate the straightforward adoption of future protocol design changes. See Fig. 4.4 for an example.

Again it should be stressed that abstraction elements are not strictly necessary for implementation, but are of great assistance for future-proof and forward-thinking design, and to take advantage of common features between protocols. Later sections will return to this topic, to discuss example protocols making use of abstraction elements, and to further demonstrate their usefulness.

---

[3] Recall the difference between general and control data-paths was discussed in Chapter 3, §3.3.0.6.

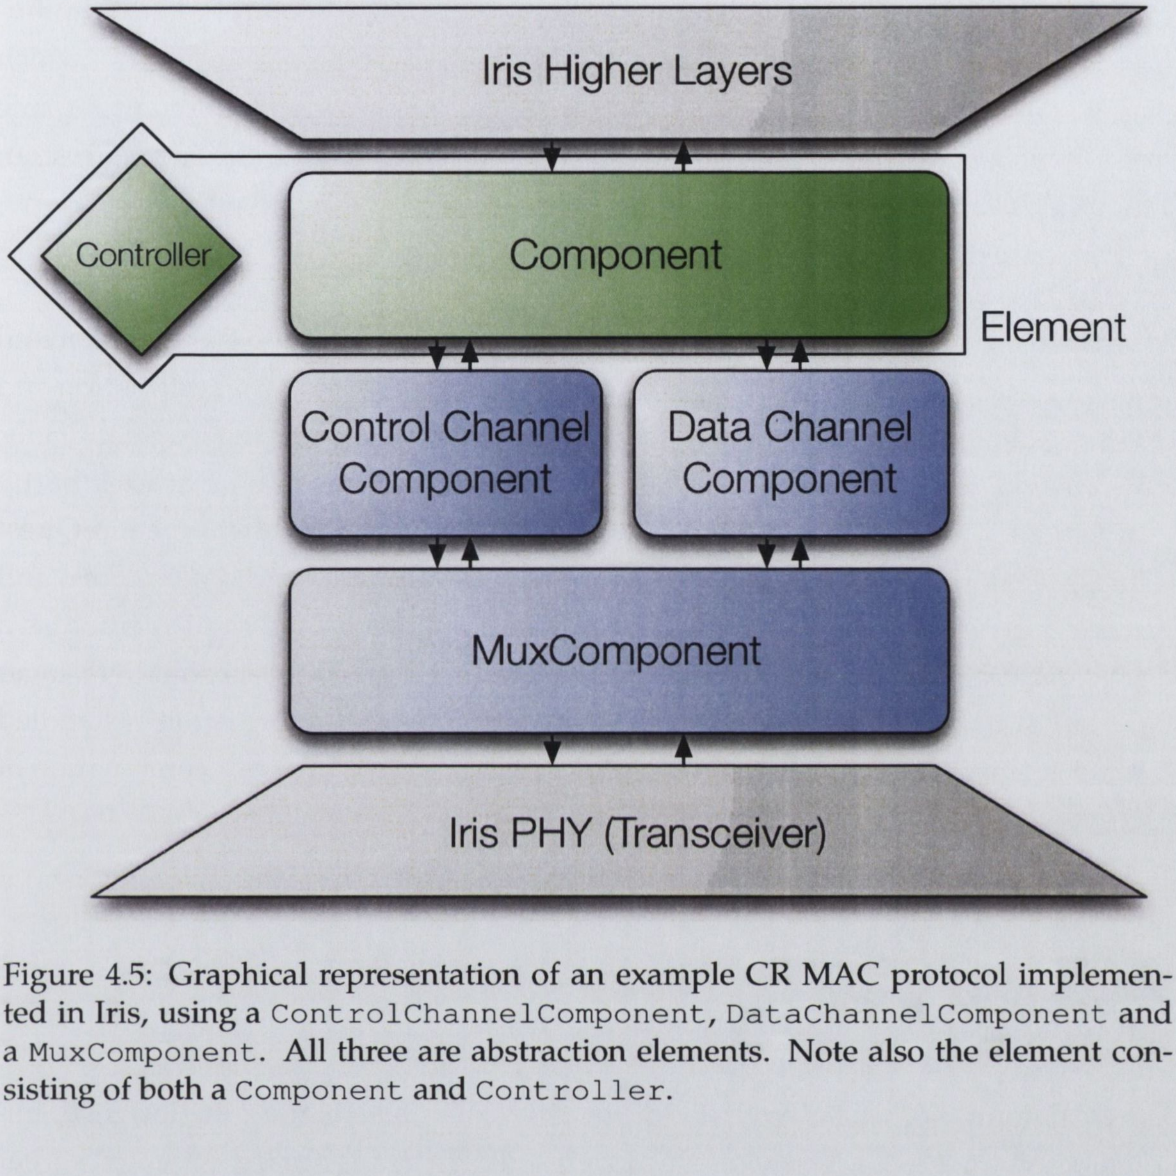

Figure 4.5: Graphical representation of an example CR MAC protocol implemented in Iris, using a `ControlChannelComponent`, `DataChannelComponent` and a `MuxComponent`. All three are abstraction elements. Note also the element consisting of both a `Component` and `Controller`.

4.1.3.3 Map Anatomy elements to Iris module(s)

There are two parts to this process. Firstly, the Iris library must be checked, to see if any required Anatomy Elements have already been implemented. The list of elements which must be newly implemented should be noted.

Secondly, each of these elements that is to be implemented must be considered in detail. As detailed in Chapter 3, the Anatomy divides a CR MAC into elements, each of which has self-contained functionality. To implement a MAC protocol using the Anatomy of a CR MAC in Iris, Anatomy elements will be implemented as one or two Iris modules. Recall that an Anatomy element can depend on access to communications data and/ or logical-information. As explained above in §4.1.1, Iris separates communication into two analogous types, data-plane and control-plane communications. Recall that modules are also divided into two types, `Components` which handle data-plane communications, and `Controllers` which handle control-plane communication. Therefore depending on the inputs and outputs it requires, an Anatomy element in a Schema will be represented in Iris as either a `Component`, a `Controller`, or both considered as a unit.

As shown in Table 4.1, communications data on the Anatomy – whether general data or control data – corresponds to the data-plane in Iris. Logical information and PHY dependencies in the Anatomy correspond to control-plane in Iris. Specifically, *if an element requires logical-information input*, then it will need to be at least partly implemented as a `Controller`. For example, a simple **Automatic Repeat reQuest (ARQ)** Anatomy element requires data input and output to send data packets and receive Acknowledgment (ACK) packet, therefore it should be implemented as a `Component`. But a more complex **ARQ** element might take into account PHY-layer error information to change its behaviour. As this logical information is handled by control-plane in Iris, it necessitate implementing the element as a combination of `Component` and `Controller`, working as a pair.

4.1.3.4 Build Iris MAC

Any Anatomy elements necessary for the CR MAC protocol must be designed, implemented, and tested if they do not already exist in the Iris library. As stated earlier in the section, an element when translated into its corresponding Iris form may consist of a `Component`, a `Controller` or both, and its implementation type will be determined by whether it relies on data-flow only, events only or both, respectively.

Anatomy Information Path:	Iris Plane:	Iris Module
Communications Data	Data-Plane	`Component`
Logical Information	Control-Plane	`Controller`

Table 4.1: Iris communications plane and Iris module types corresponding to Anatomy information path categories. If both information path categories are required by an element, then it must be implemented in Iris as a pair of both modules. (Refer to Chapter 3, Fig. 3.5 for a breakdown of these.)

4.1.3.5 Build Iris PHY

PHY layer DSP blocks must be combined to build a PHY satisfying the listed dependencies, and meeting the contract. The existing library of `Components` within Iris may be combined to produce the transceiver(s) necessary, or new PHY `Components` may be constructed. Provided the contract is accurate and all dependencies are clearly noted, any PHY(s) satisfying them should be usable by the MAC protocol.

Any assumptions or features offered by hardware that a protocol wishes to take advantage of in its design, must be represented as a contract. In the case of Iris, the PHY layer OFDM modulation and demodulation components, as modified for MAC protocol experimentation, offer two guarantees:

- Any packets passed to components above them will have passed a checksum
- Packets boundaries will have been determined, so that blocks of bytes received by the MAC correspond precisely to complete packets.

Therefore these behaviours can be represented as a PHY contract on the Anatomy, by any CR MAC protocol wishing to take advantage of them in their design.

4.1.3.6 Create Iris xml

Bearing in mind the extra `Components` introduced in §4.1.3.2, an Iris xml must be constructed to indicate the protocol's required Iris modules, both MAC-layer and PHY-layer, and to specify the links between them. This is little more than a linear representation in xml of the diagrammatic version of the radio, as was shown earlier in Fig. 4.1. As an example, the structure of the xml configuration that would correspond to Fig 4.5 is outlined below, to demonstrate how it captures the radio's structure.

```xml
<softwareradio name="SimpleRadio">
  <controller class="greencontroller"/>
  <engine name="simpleradioengine" class="stackengine">

    <component name="higher_layers" class="higherlayers">
      <port name="down" class="io"/>
    </component>
    <component name="greencomponent" class="somecomponent"/>
      <parameter name="parameter1" value="somevalue"/>
      <parameter name="parameter2" value="someothervalue"/>
      <port name="up" class="io"/>
      <port name="downcont" class="io"/>
      <port name="downdata" class="io"/>
    </component>
```

```
    <component name="cont_chan" class="abstractcomponent"/>
      <port name="up" class="io"/>
      <port name="down" class="io"/>
    </component>
    <component name="data_chan" class="abstractcomponent"/>
      <port name="up" class="io"/>
      <port name="down" class="io"/>
    </component>
    <component name="muxcomponent" class="muxcomponent"/>
      <port name="upcont" class="io"/>
      <port name="updata" class="io"/>
      <port name="down" class="io"/>
    </component>
    <component name="iris_phy" class="iris_phy"/>
      <port name="up" class="io"/>
    </component>
  </engine>
</softwareradio>

<link source="higher_layers.down" sink="greencomponent.up"/>
<link source="greencomponent.downcont" sink="cont_chan.up"/>
<link source="greencomponent.downdata" sink="data_chan.up"/>
<link source="cont_chan.down" sink="muxcomponent.upcont"/>
<link source="data_chan.down" sink="muxcomponent.updata"/>
<link source="muxcomponent.down" sink="iris_phy.up"/>
```

Each `Component` requires its own `<component>` block. A `<component>` block begins with a unique `name` by which each `Component` in a radio is distinguished. Then the `class`, which is the variety of `Component`, is specified. A radio may instantiate many instances of the same class of `Component` simultaneously. Note that in this example, there are two instances of "`abstractcomponent`", but each has a unique name. Depending on the corresponding C++ implementation of a `Component`, it may have one or more parameters, optionally set at build-time using the `<parameter>` tags. Note the "`greencomponent`" has two example parameters. Each `Component` also has one or more ports, through which communications data passes in and/ or out. These are uniquely named for a `Component` in its C++ implementation. Any `Controllers` are included at the top of the xml. Note that a `Controller` does not receive a unique name, and therefore only one from each class can be used in a radio. Finally, radio structure is specified by building a directed graph of arcs between `Component` ports. Each `<link source=?? sink=??/>` tag at the bottom of the xml, creates a directed arc, by specifying a component source and sink port by name. For a more detailed discussion of Iris xml configurations and options, the reader is referred to explanatory works by Lotze and Sutton [79, 83].

A full Iris xml representation of the radio later discussed in this chapter can be found in Appendix A.

4.1.3.7 Run

Once the Iris xml is complete, it can be passed to the Iris core and run for testing and/ or operation.

4.2 CycloMAC: Implementation of a CR MAC in Iris

To showcase the Anatomy in action for a CR MAC, this section does the following: First, it presents a scenario requiring a CR MAC, before using the Anatomy to design a protocol suited for that scenario. Then, it details the implementation of this protocol in Iris using the Anatomy, illustrating the method introduced in §4.1.3 by following through it step-by-step. This implementation is a precursor to the next chapter, which will discuss the benefits the Anatomy brought to the design and the implementation process.

For reasons that will become clear later in this section, the newly designed protocol will be referred to as "CycloMAC".

4.2.1 Scenario

Recall that the most commonly considered motivator for a CR MAC in the literature is Dynamic Spectrum Access (DSA)[4]. In particular opportunistic DSA is so commonly the motivating scenario of choice that it is appears to be synonymous with CR to many authors. Though this constant conflation is misleading, it is not necessarily an indictment of the challenges involved in realising opportunistic DSA, at the MAC layer or any other. In fact, given an appropriately challenging set of restrictions to the scenario, opportunistic DSA becomes an excellent test of the capabilities of a CR.

In light of the above, and to frame it in terms comparable with protocols proposed in the literature, this section will design and implement a CR MAC protocol using the Anatomy for an opportunistic DSA scenario. As discussed in Chapter 2, §2.1.2.1, this means that devices employing this MAC protocol will opportunistically take advantage of spectrum licensed to others, transmitting in it when it is idle. As already stated, with an appropriate set of restrictions imposed by a scenario, opportunistic DSA is a strong challenge for a CR. A lack of stringency in the initial

---

[4] Dynamic Spectrum Access (DSA) was introduced in Chapter 2, §2.1.2.1.

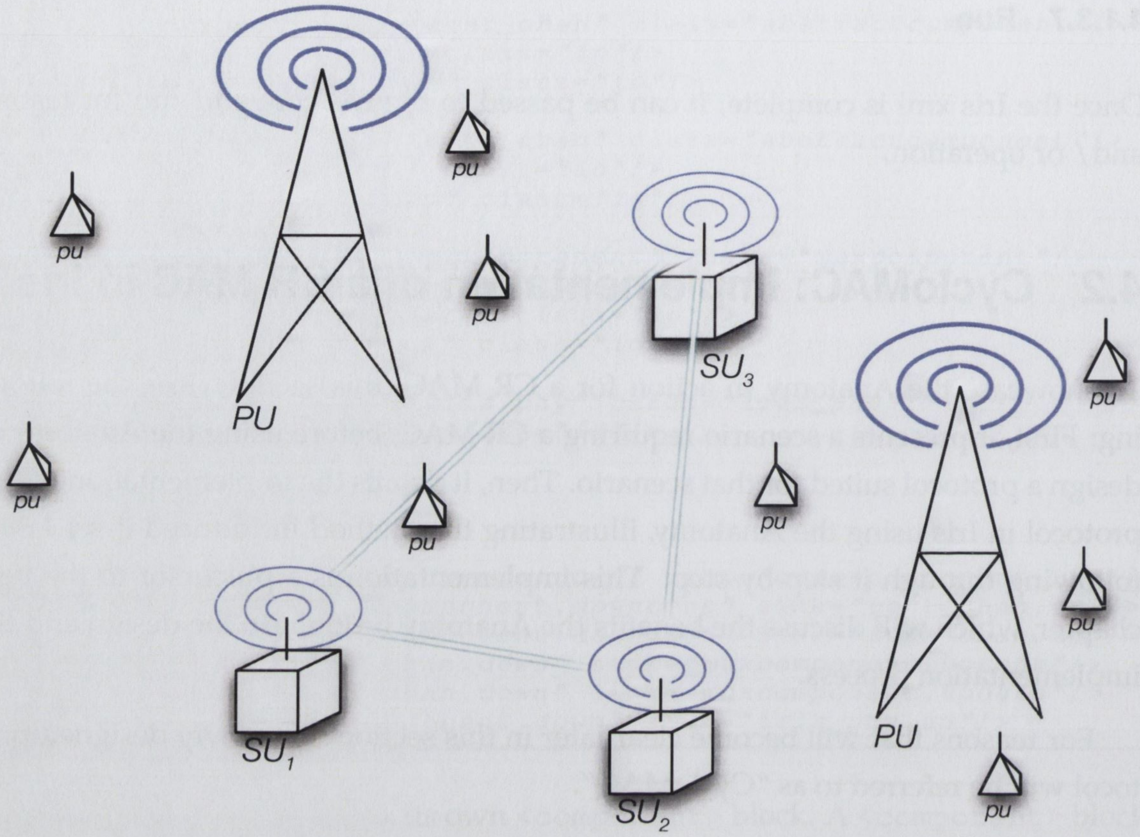

Figure 4.6: Diagrammatic representation of this scenario, depicting three SU devices, labelled SU_1-SU_3, and a large number of PU base-stations and terminal devices. As all devices are geographically close, transmission on overlapping frequencies by multiple devices may cause destructive interference.

scenario restrictions imposed by many authors is a failing of many protocol proposals in the literature, as it may result in inflexible, uninnovative and unimplementable protocols. This section will pose a scenario which is sufficiently challenging to avoid such a failure, and demonstrate that even provided with such a stringent set of requirements, the Anatomy can be used to design an appropriate CR MAC solution in a clear and rapid manner.

A diagrammatic representation of the scenario is shown in Fig. 4.6, depicting 3 Secondary User (SU) devices and a number of Primary User (PU) devices.

In this scenario, no dedicated spectrum will be assigned to CR devices, nor will any part of the spectrum be guaranteed to be consistently unused by PUs. The effect of this will be to restrict the CR MAC protocol's ability to rely on constant availability of a specific frequency to use for a static Common Control Channel (CCC) for signalling. More importantly, it will mean that such a frequency cannot be used for initial setup. In short, at switch-on, devices have no conventional way to ascertain the state of the network. How many, if any, other devices are active and what frequency or frequencies they are operating at will be unknown. While a significant restriction, particularly when conventional MAC techniques are considered, it is a realistic one to consider for an opportunistic DSA scenario. PU priority is likely to be absolute, and therefore an SU cannot under any circumstances be dependent on the availability of a specific piece of spectrum.

To reduce deployment cost and reliance on some devices over others, construction of device-supporting infrastructure will not be permitted. The effect will be to force the protocol to use distributed techniques. No one device employing this protocol may be more important than others, or in some way critical to the functioning of the overall system. Though this a somewhat subjective design decision, as CR and DSA techniques do not explicitly mandate networks be centralised or decentralised in structure, there is an intuitive appeal to marrying a CR with the resilience and redundancy of a distributed protocol. Furthermore, the redundant nature of such a system and lack of a single coordinating master device makes it more challenging than a centralised approach.

To reduce cost and PHY complexity, devices will be limited to a single PHY transceiver each. The effect of this will be to confine all control signalling to being handled by the single transceiver. Moreover, it will limit the device to transmitting and receiving on a single contiguous block of spectrum, unless it makes use of unconventional PHY techniques. Though many modern RF communications devices contain multiple discrete transceivers, each adds additional cost to a device. The highly flexible PHY transceivers employed in a CR are particularly costly, and it benefits a device to make use of protocols requiring a minimal number of such transceivers.

In terms of the capabilities that may be reasonably required of the PHY in this scenario, there will be no strong restriction. However, any unusual requirements must not require significant or potentially costly modifications. Therefore it should use a modulation scheme and structure that is conventional or close to it, and any DSP required at the receiver should not be a dominant factor in the overall computational complexity of the device.

The majority of CR MAC protocols as summarised in Chapter 2, §2.4.1, assume that frequency has been divided into discrete orthogonal channels. The origin of this channel map is usually not discussed, but it is implicitly assumed to be known and used both by PU and SU devices in opportunistic DSA protocols. For this scenario, a discrete channel map is not available to or used by any device, PU or SU. The effect of this will be to increase the difficulty of locating other CR devices, as the list of possible centre frequencies will not be fixed. Also, as the lack of channel map applies to PUs as well, it will prevent anticipation of future PU centre frequencies and presumption of their having a fixed transmission bandwidth. This restriction is particularly realistic, at least in so far as it separates the possible centre frequencies of PUs from SUs. It is also of particular challenge, in the difficulty it adds to locating other devices, both PU and SU.

The final requirement of the scenario is one common to most opportunistic DSA MAC protocols – the requirement to avoid use of any spectrum that overlaps with PU transmissions. In order to do this, a CR must be in some way aware of the activity of PU devices. Additionally, it must have mechanisms in place to allow it to cease its own communications, change frequency, and recover communications in the event of an unexpected PU arrival.

This scenario is a reasonable emulation of the TV white-spaces access rules for which are beginning to come in to force in several regulatory jurisdictions. The Federal Communications Commission (FCC) in the USA [90] and Ofcom in the UK [91] have both issued such rules. The scenario could also be considered analogous to any model where a regulator has permitted secondary access to spectrum by CR devices. As regards TV white-spaces access, a process of consultation remains ongoing as to the exact restrictions that will be posed on opportunistic DSA devices operating in these bands in many countries. Nonetheless, the set of scenario restrictions outlined here can be expected to be similar to the sorts of challenges a CR MAC protocol for TV bands will be expected to face across the world, as legal requirements continue to be published. In fact, the hardware restrictions and lack of spectrum channelization in this scenario are likely more stringent scenario restrictions than might be required of TV white-space devices in the real world.

4.2.2 Ramifications of the scenario

The scenario restrictions outlined in the section above will have many ramifications on a potential CR MAC protocol being tailored to them. Though some of these ramifications were discussed in general terms above, this section will consider them in more specific terms, with the aid of an Anatomy Schema. These scenario restrictions are enumerated and summarised in Table 4.7a, labelled S1-S5.

Fig. 4.7b lays out a simplified Anatomy Schema, with no elements detailed as

Scenario restriction	Effect on implementation
S1. No dedicated spectrum available or guaranteed idle	Devices cannot bootstrap in pre-defined frequency; Hardware cannot be frequency limited to reduce cost; No CCC
S2. No centralisation or heterogeneous device capability	Protocol must be decentralised; All devices must be capable of performing all roles in the protocol.
S3. Single PHY transceiver	Device may not monitor multiple frequencies; Constantly active CCC not viable.
S4. No *a priori* channel map, used by SU or PU	Bootstrap/ rendezvous cannot use a channel map to restrict scope of search; Orthogonality with a PU can never be assumed.
S5. Avoid PU occupied frequencies	Awareness of PU activity needed; Mechanism to suspend communications, change frequency and recover required.

(a) List of restrictions imposed by the scenario, labelled S1-S5, and their corresponding effect(s) on any implemented CR MAC protocol to be designed for it.

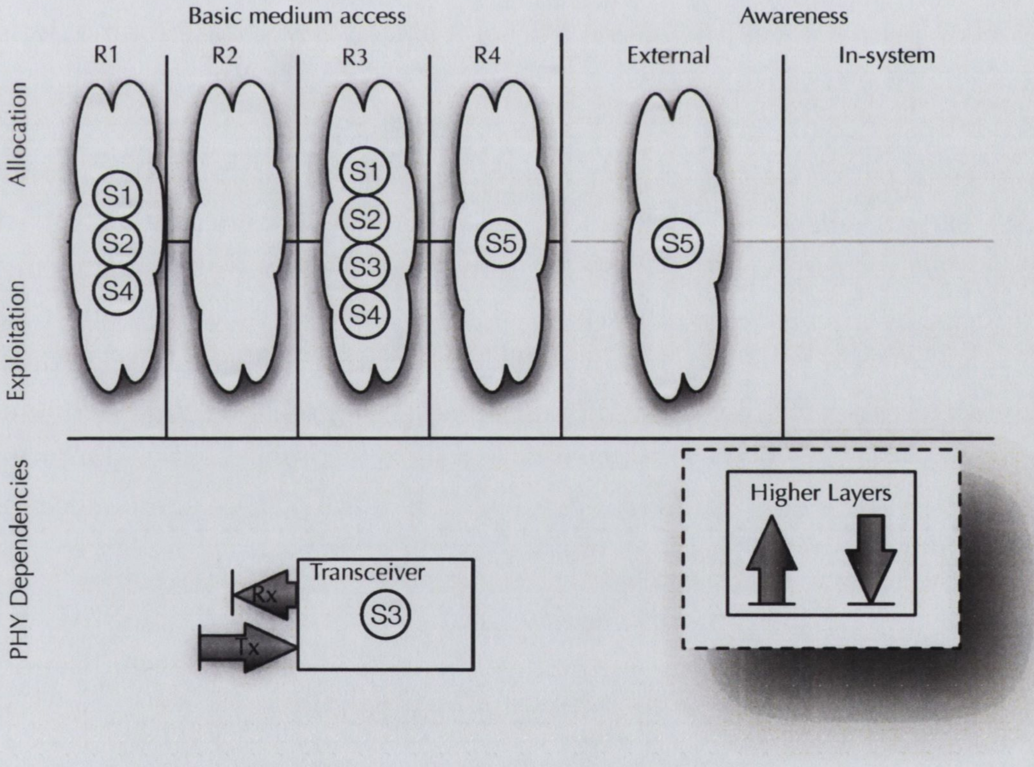

(b) The ramifications of the scenario on a potential CR MAC protocol, laid out on an Anatomy Schema.

Figure 4.7: Table of scenario restrictions, and a simplified Anatomy Schema noting their effects on a potential CR MAC protocol.

yet. Recall the four responsibilities of a CR MAC which are here relabelled R1-R4, as follows:

- **R1** ***Define***: It must choose when and where to access the medium, as well as it what it is.
- **R2** ***Deliver***: It must ensure message delivery.
- **R3** ***Control & Coordinate***: It must arranging any necessary control signalling, and if necessary the means for the signalling to take place.
- **R4** ***Edge-case***: It must handle any edge cases.

The simplified Schema makes use of the scenario restriction numbering imposed by the table overleaf, considering their effects on each of a MAC's four responsibilities, and on its awareness. Traversing the Schema left-to-right, each column is now considered in turn.

Responsibility R1 of a potential CR MAC protocol is affected by several of the restrictions. Restriction S1, that no dedicated spectrum is available or guaranteed idle, affects it primarily by preventing it from making a static design-time decision as to where in the spectrum resource it accesses. Instead it must be flexible, take into account its current situation, and not make assumptions. In particular, it cannot limit the potential range of frequencies to be accessed, as there is no guarantee the chosen subset will ever be available for use. It will also be affected by restriction S2, in that it cannot take a centralised approach to this decision. Whatever means by which it decides when, where and what spectrum to access must be distributed in its implementation. S4 affects a potential CR MAC by requiring the MAC to deal explicitly with frequency choice from an infinite range of possible centre frequencies. In disallowing the use of a set of discrete frequency channels, either the protocol must explicitly channelise the frequency itself, or it must have some mechanism that allows it to avoid the necessity of doing so.

Responsibility R2 does not provide any specific challenges for a CR MAC in this scenario. Provided all other responsibilities are sufficiently addressed, a conventional mechanism for ensuring delivery may be employed.

Responsibility R3, a CR MAC's responsibility to control and coordinate communications, is strongly affected by the scenario's restrictions. Scenario restriction S1 means that a fixed dedicated CCC cannot be employed for control. Furthermore, it means that at device start-up, no channel will be available in order to allow initial coordination. A means of initial coordination that is not dependent on any specific frequency must be employed. Restriction S2 will have some effect on this responsibility also, in that no device may be indispensable to handling control or coordination. Any mechanisms employed must be fully distributed. Restriction S3 means that a transceiver cannot be dedicated wholly to whatever control and coordination mechanism is employed, nor can any method requiring more than one transceiver be used. And restriction S4 affects coordination, by negating the possibility of using a channel map to narrow the range of possible center frequencies that must be "searched" in order to determine if other CR devices are active. Instead there is an infinite range of possible center frequencies.

Meeting responsibility R4 in a potential CR MAC for this scenario will depend primarily on the elements chosen to address the other responsibilities. If not addressed in other elements however, there will be a requirement for some mechanism relating to scenario restriction S5. That is, the edge case of a PU return to a channel must be fully provided for. A mechanism to suspend communications, arrange a change of frequency, and restart communications will be necessary to fully handle this.

Finally, external awareness is affected by scenario restriction S5 – any potential CR MAC protocol for this scenario must be aware of current PU activity.

Also, restriction S3 is noted on the transceiver in the Anatomy Schema, to indicate that the transceiver count is limited to one.

In order to summarise consideration of a potential CR MAC for, it is enlightening to now consider on a device level some example interactions that the CR MAC protocol will be required to facilitate.

4.2.2.1 Device switch-on

Consider a CR device. When switched on, there are two possible situations the device may find itself in: (a) It is the only active CR device or (b) Other CR devices are active. The device's first requirement is the ability to determine which situation it is in, (a) or (b).

If it is in situation (a), it will then require some procedure it can employ in order to allow it to inform devices activated in future that *they* are not the only active CR device. This procedure must take into account PU activities so as not to disrupt them.

If it is in situation (b), it will need a means to find the centre frequency and transmission parameters of the other devices. And once it has done so, it will need an agreed handshaking mechanism to allow it to commence communication with them.

4.2.2.2 Device switch-off

Consider that same CR device, now in communication with several other devices. In order to fully meet the scenario restriction for decentralisation, the removal of the device – or indeed many CR devices – should not cause remaining devices to lose the ability to communicate.

4.2.2.3 PU arrival on overlapping centre frequency

Finally, a third example interaction. Consider a group of CR devices, all communicating on some frequency f_a. If a PU commences transmission on any frequency overlapping with f_a, group communication must stop as quickly as possible. By some means, the group must relocate to a new frequency f_b, which does not overlap with any PU transmissions. A decision will need to be made as to what centre frequency to choose as f_b, and all devices will require a means to find this value. The end result of such an interaction should be all devices fully resuming communication, on this new frequency f_b.

4.2.3 Addressing the scenario

The purpose of this section is to create a more detailed Anatomy Schema, by considering how a CR MAC suited to this scenario might operate. It will revisit the responsibilities of a potential CR MAC, and outline the tentative placement of Anatomy elements for a protocol. It will also propose candidate methods to be used by these elements where relevant.

In attempting to create a more detailed Anatomy Schema, one of the first considerations must be the role(s) the CR MAC will play. As device-to-device communication must be facilitated, it is clear the protocol will need to play an exploitation role. But though not explicitly required by the scenario, more flexible behaviour and scalability will be allowed for if it also plays an allocation role. Therefore in order that the CR MAC being implemented be as fully featured as possible, it will play both a spectrum resource allocation and exploitation role. Note that the scenario requires the protocol to be distributed in nature. Therefore, in playing its allocation role it must be either transmit-receive pairs or groups of devices that are being allocated to.

Each responsibility column of the Anatomy Schema will now be revisited, in order to establish tentative elements for a protocol suited this scenario. This step relies on a CR MAC designer's knowledge of a range of candidate protocol elements and their dependencies, so that they may be employed while not violating the scenario restrictions. Over time and by examining the literature, a catalogue of such protocol elements can be built up by a designer. When doing so it is important that for every element its strengths and weaknesses be noted. In particular, each element must be noted for the responsibility or responsibilities it meets, and its dependencies. These dependencies in turn imply the restrictions it can comply with. In this section, such a knowledge of a sufficient number of appropriate candidate elements is assumed.

4.2.3.1 R1 Define: Tentative elements

In order to meet responsibility R1 in playing its allocation role, a "frequency selection" element will be required. It cannot be designed to make a default choice, and will instead need to make a decision taking into account current PU activity. As will be returned to later, this indicates a likely logical-information input requirement. At the exploitation level, as only a single transceiver is available, the choice of Time Division Multiple Access (TDMA) to divide the spectrum resource allocated by the "frequency selection" element would be a reasonable one. This will mean that there will be no further subdivision in frequency of the spectrum resource allocated by the "frequency selection". However, scenario restriction S2 means that all techniques must be distributed in nature. This will prevent any one device from being a permanent master in the TDMA system, and require a decentralised TDMA implementation be employed. Note that the decision to employ TDMA will also assist in meeting responsibility R3, as explained below.

4.2.3.2 R2 Deliver: Tentative elements

To meet responsibility R2, the protocol will require a means to ensure delivery. Provided all other responsibilities have been met, ARQ is a suitable candidate for this.

4.2.3.3 R3 Control & Coordinate: Tentative elements

Responsibility R3 may require several elements to be fully met. In the context of its allocation role, and as discussed earlier, CR devices will require a mechanism for coordination at start-up, which is not frequency dependent and does not rely on channellisation. This thesis will refer to this problem as *bootstrap*. A tentative element will likely be required to handle it. Also as part of responsibility R3, is the CR devices' need to be capable of locating one another across a potentially wide range of frequencies. After all, the decision of where to access the medium is pointless if no other device you wish to communicate with is aware of it. This problem is referred to as *rendezvous*[5]. Rendezvous is also considered to be a tentative element, though depending on the mechanism chosen, it may be strongly tied to the bootstrap element.

In the context of this potential CR MAC, there may have been question as to which row of the Anatomy these tentative elements should be placed on; which role they play a part of. It is logical that they be indicated to play a part of the spectrum resource allocation role, as they both play a part in allowing devices be made aware of the decision of the "frequency selection" element indicated as part of column R1 earlier.

Several solutions have been proposed in the literature to solve the bootstrap and rendezvous problems, usually involving channelising the available frequencies and employing carefully chosen sequences of channels on which to alternate between listening for and transmitting to other CRs devices [93]. However, particularly bearing in mind the restrictions of the scenario, a promising candidate solution to both problems is to employ *cyclostationary signatures*.

While cyclostationary *features* are a property of the periodicity of all manmade RF communications, a cyclostationary signature is a feature that can be intentionally embedded in a waveform [94, 95]. By embedding a signature in a waveform at a pre-agreed cyclic-frequency, a transmitter makes its signal easy to pick-out for a receiver sweeping across a large range of frequencies. The presence – or absence – of a signature can be detected using a computationally lightweight process that does not require demodulation of the signal. They are robust in the presence of noise,

---

[5] This can sometimes be referred to as "neighbour discovery" [92]. Recall that §2.2.3 introduced the *multi-channel rendezvous* problem in the context of multi-channel MACs. Rendezvous here is a more general case of the multi-channel rendezvous problem, where not just pre-defined channels, but any possible frequency in range could potentially be home to a device with which we may wish to communicate.

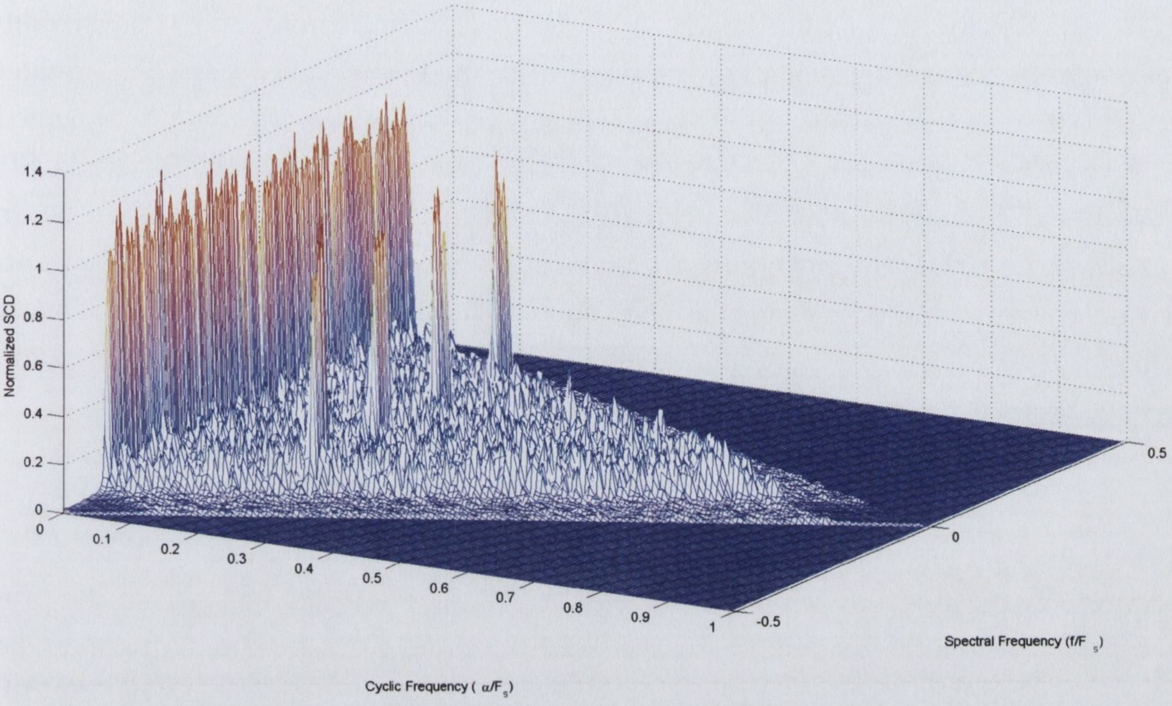

Figure 4.8: Capture of a signal with four intentionally embedded cyclostationary signatures. Each peak at cyclic frequency 0.2 is a signature.

and detectable even when embedded in intermittent signals, a behaviour crucial for their use in a CR MAC. These signatures are easily manipulable, in that it is possible to generate and distinguish between a range of unique signatures. Fig. 4.8 reproduces the two-dimensional spectral correlation function of a captured signal from [95]. It has had four cyclostationary signatures intentionally embedded in it, at a cyclic frequency of 0.2. Therefore, a specific cyclostationary signature can be used as a unique identifier of a device, or group of devices. Once a signature is detected, the calculation of the exact centre-frequency at which the waveform containing it is being transmitted is trivial. A CR MAC attached to a PHY capable of transmitting and detecting cyclostationary signatures can make use of this detection capability to locate other transmitting devices in frequency, or transparently inform searching devices of its own frequency of operation while engaging in ordinary transmissions.

Recalling that the scenario outline in §4.2.1 presented the general limitation that PHY dependencies not require significant or costly modifications, it should be noted that cyclostationary signature embedding and detection is computationally lightweight and straightforward to implement. At the transmit side there is effectively no processing overhead. The only requirement is that a multi-carrier modulation technique such as OFDM be used[6] , and that data must be mapped to be transmitted simultaneously on a small group of pairs of pre-agreed subcarriers. This creates a consistent spectral-correlation peak. Detection and location of this pre-arranged peak at a receiver requires straightforward Fast Fourier Transfrom (FFT)-based DSP. Such FFT DSP capability is usually available effectively "for-free" in many modern multi-purpose communications hardware solutions, or can be implemented in software using mature processing libraries such as the FFTW C libraries [96]. Clearly therefore, cyclostationary signatures offer a solution to the scenario's bootstrap and rendezvous problems that complies with the PHY complexity limitations.

Use of a cyclostationary signature also removes the issues associated with a lack of channelisation. When rendezvous is performed using the signatures, the precise centre frequency of the signal being searched for is determined as part of the process and need not be known in advance. Furthermore, embedded signatures can be further extended to allow a receiver to determine a transmission's signal bandwidth without a priori information [97]. Use of signatures will allow the protocol to meet CR MAC responsibility R3, without violating scenario restrictions S1, S2 and S4.

It should be noted that existing work using cyclostationary signatures performed at CTVR was solely on a point-to-point basis, involving single transmit-receive pairs of devices. A contribution of this thesis is to establish by experimentation (i) that intermittent signals with embedded signatures can still be reliably detected and (ii) the transmission of the same signature by multiple devices does not hinder the ability of other devices to reliably detect it.

Finally, third element required in column R3 will be one to be responsible for in-system control message passing as part of the protocol's exploitation role. The lack of availability of a consistent CCC, combined with scenario restriction S3 which mandates only a single PHY transceiver be used, implies either a need for a single transceiver based hopping solution to the control problem, or the use of split-phase (in-band) control. Given the choice of TDMA in column R1, it is reasonable to make use of a split-phase control mechanism, in which part of each time-slot is reserved for control message exchange. A tentative element to represent this "in-band control channel" will therefore be indicated on the Anatomy Schema.

---

[6] Recall that the OFDM `Component` in Iris was already modified to perform well for MAC experimentation, as mentioned in §4.1.2.

4.2.3.4 R4 Edge-cases: Tentative elements

Edge-cases are caused primarily by elements chosen to handle responsibilities R1-R3 and awareness. As the tentative elements listed thus far do not suffer any notable issues, the only edge-case requiring handling is the event of a PU arrival on frequency overlapping with current transmission. As frequency in this MAC protocol is selected as part of the allocation role, a tentative "Handle PU returns" element is indicated in the allocation row.

4.2.3.5 External Awareness: Tentative elements

Avoiding overlap with PU occupied frequency is extremely important in an opportunistic DSA scenario. In order to do this, a source of information on PU activity is required to inform any frequency selections made. A tentative element will therefore be marked as"PU activity info source". As frequency selection in this protocol

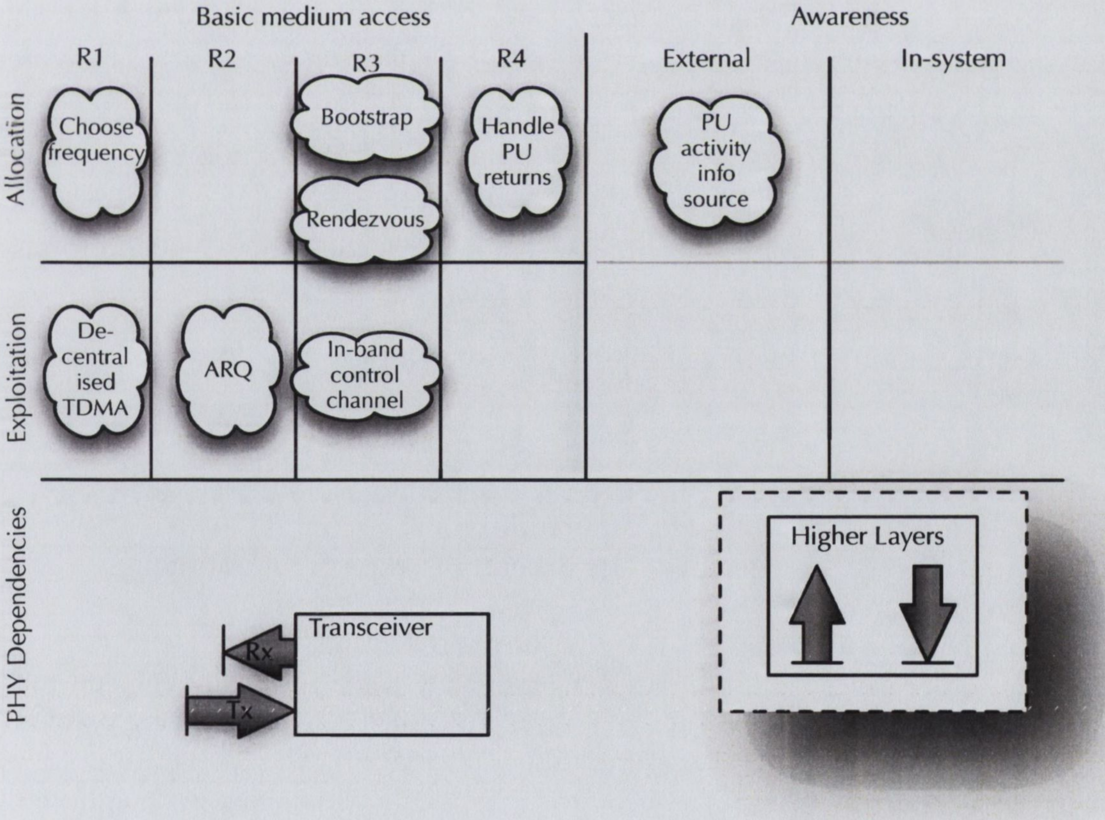

Figure 4.9: Tentative Anatomy elements for a CR MAC protocol designed to meet this scenario.

is performed as part of its allocation role, this element will also be marked in the allocation row.

Discerning instantaneous availability of a a given swathe of spectrum at a given point in time is often the responsibility of spectrum sensing. There is great variety in the range of spectrum sensing technologies that can be employed, including Energy Detection, Matched Filtering, Wavelet detection, detection tailored specifically to known properties of a PU, and cyclostationary feature detection [8]. Each technology comes with benefits and drawbacks that must be taken into account in choosing which to deploy.

Cyclostationary feature detection as mentioned above should not be confused with cyclostationary *signatures*, which were proposed earlier as a candidate mechanism to assist in meeting responsibility R3. Cyclostationary features are inherent artefacts of a waveform, while cyclostationary signatures are additional features that have been intentionally embedded. Detection of intentionally embedded cyclostationary signatures is a less computationally challenging task than feature detection, as detections only take place on the pre-agreed cyclic frequency at which the signatures are to be found. Feature detection usually requires a sweep across a wider range of possible cyclic frequencies, which requires more time and processing resources. It is worth noting that were the limitation on computationally complex PHY requirements mentioned in §4.2.1 to be sufficiently relaxed, the same modifications involved in detecting cyclostationary signatures could be scaled up to also provide feature detection based environmental awareness.

The most commonly considered alternative is to employ a geolocation database.

Scenario restriction	Chosen solution
S1. No dedicated spectrum available or guaranteed idle	Use cyclostationary signatures for bootstrap and rendezvous.
S2. No centralisation or heterogeneous device capability	Use decentralised protocol techniques.
S3. Single PHY transceiver	Use split-phase control.
S4. No *a priori* channel map, used by SU or PU	Cyclostationary signatures remove need for pre-agreed channel map by allowing precise frequency calculation.
S5. Avoid PU occupied frequencies	Refer to geolocation database when choosing frequency of operation.

Table 4.2: List of restrictions imposed by the scenario S1-S5, and the solutions chosen to address them for this implementation.

This is a database that records the location, activity and transmission parameters of PU devices, maintained by a third-party such as a government regulator. A CR device with knowledge of its own location can access and monitor this database to determine PU whereabouts. Real-world DSA use-cases being considered are currently leaning towards such databases as the spectrum occupancy information source of choice [98, 99]. As a result, for this scenario the availability of a geolocation database is assumed. Furthermore, the details of the method of access of the database, and the means necessary to obtain the CRs' location will be considered beyond the scope of this implementation.

An Anatomy Schema representing all tentative Anatomy elements discussed in this section is shown in Fig. 4.9. Table 4.2 summarises the candidate solutions discussed in this section, as they relate to each scenario restriction.

4.2.4 Implementation Details

This section will now discuss the details of implementation of the proposed protocol. It will follow the procedure for implementing a CR MAC protocol using the Anatomy in Iris as explained in §4.1.3. In doing so it will flesh-out the details of the tentative elements proposed to address scenario restrictions in the previous section.

4.2.4.1 Create Anatomy Schema

Step 1 of the procedure is to create an Anatomy Schema. As an Anatomy Schema indicating tentative elements has already been laid down, this step will involve confirming the elements as they will appear in the final protocol.

Traversing the Anatomy Schema in Fig. 4.9 left-to-right, first consider only its exploitation row.

The proposed CR MAC protocol requires a decentralised TDMA implementation. Furthermore, sub-frames within the TDMA frame structure will be dedicated control phases. Though it is represented by a single tentative element in Fig. 4.9, it is worth briefly considering whether there might be a benefit to implementing it as more than one element. As a goal of the Anatomy is to enable reuse, it is worthwhile throughout the design process to attempt to make design choices that facilitate element reuse at a later date. A TDMA mechanism could be considered to consist of two parts: one responsible for determining when exactly frames start (its synchronisation mechanism) and the other for actually scheduling data into the frames for transmissions (its scheduler). If functionality were to be split along these lines, one might easily replace the synchronisation mechanism without needing to affect the scheduler. Therefore here functionality will be split over two such elements, a scheduler and a synchronisation element. Specifically, a **Decentralised synchronisation** element will be implemented, along with a **Scheduler** potentially agnostic as to whether they system is centralised or de-centralised.

Moving on, if ARQ is employed to ensure delivery, an element modelled on the description of ARQ from Chapter 3, 3.3.0.5, can be implemented and used. Provided it is implemented to make use of abstraction elements to split data and control messages, it will be agnostic of the fact that it is effectively sending its ACK packets in the control phase, while sending data packets in the main data phase of each frame.

Finally, an element will be required to indicate the **In-band control channel**, receiving control packet input from the **ARQ** element, and routing it to the control subframe input on the **Scheduler**.

Considering then the allocation row of the protocol, several tentative elements have been proposed to handle bootstrap, rendezvous, frequency choice and PU returns. As per the suggestion in the previous section, the protocol will employ cyclostationary signatures to allow device detection and identification. These can serve as part of both the rendezvous and bootstrap process. Furthermore, considering the choice of frequency, the two situations in which this choice must be made are
i.) as part of the bootstrap process, in the case where no other device is currently active and ii.) when an alternative frequency choice is required owing to a PU return. Given the interrelationship between all these tentative elements, it is reasonable to incorporate their responsibilities as one single **Cyclostationary bootstrap, rendezvous & recovery** element. To indicate its role in handling all of these responsibilities, it will be depicted crossing all of the relevant columns.

It is important to note that the embedding and detection of cyclostationary signatures is a physical-layer process. Therefore this element will be responsible for the control and interpretation of signatures, but not for the actual DSP involved in incorporating the signature into the waveform, and the detection of it.

To complete consideration of the allocation row then, is the need for a PU activity information source. As discussed above, the existence of a geolocation database will be assumed for this implementation. In the context of the Anatomy Schema for this protocol, this can be represented as a single **Geolocation DB** element.

Though these will be reinforced fully as part of the later steps of the implementation process, it is important to consider for a moment the logical-information requirements between elements. In order to indicate the start of frames to the **Scheduler** element, the **Decentralised synchronisation** element will need to provide it as a logical-information output. Additionally, as will become clear later in this section, in order to fully effect a recovery mechanism in the event of PU arrival on the currently occupied frequency, the **Cyclostationary bootstrap, recovery & rendezvous** element will require knowledge of which device is currently acting as lead in the group of active transmitting-devices. This information will need to be output by the Decentralised synchronisation element. The final logical information requirement is the need for knowledge of current PU activity at the **Cyclostationary bootstrap, recovery & rendezvous** element, so that appropriate frequency choices can be made. This necessitates a logical-information input from the **Geolocation DB** element.

The proposed CR MAC protocol is sufficiently complex that it has a multitude of PHY contracts and dependencies, some of which have already been hinted at. This protocol will require: a cyclostationary signature compatible PHY transceiver[7]. It will require an additional meter in the form of an output of cyclostationary signature status[8]. It will require knobs giving the ability to change current frequency of operation and the ability to enable and disable communications data-passing to and from the MAC layer. This requirement is needed in order to allow the **Cyclostationary bootstrap, rendezvous & recovery** element to disable the MAC when the PHY transceiver is performing a sweep of frequencies to detect a cyclostationary signature. All of these contracts, knobs and meters must all be clearly marked on the Anatomy Schema. Furthermore, recall that earlier, in §4.1.2, modifications were made to Iris's OFDM `Component`, to allow it perform better when driven by MAC layer. These modifications mean that packet boundaries are automatically determined, and packets passed on by the PHY are guaranteed error free. These can be marked as contract on an Anatomy Schema. Finally, it should be noted that any TDMA MAC requires finely-grained control of time, and therefore this requirement should also be marked as a PHY contract, and the presence of errors in attempt to schedule a packet should be fed-back via a meter.

Gathering these decisions, the Anatomy Schema of the proposed CR MAC protocol is depicted in Fig. 4.10, dubbed "CycloMAC" owing to its significant use of cyclostationary signatures.

---

[7] This is marked as contract *Cyc-sig active* on the Anatomy Schema shown overleaf, in Fig. 4.10.

[8] This is marked as meter *Cyc-sig status* on the Anatomy Schema shown overleaf, in Fig. 4.10.

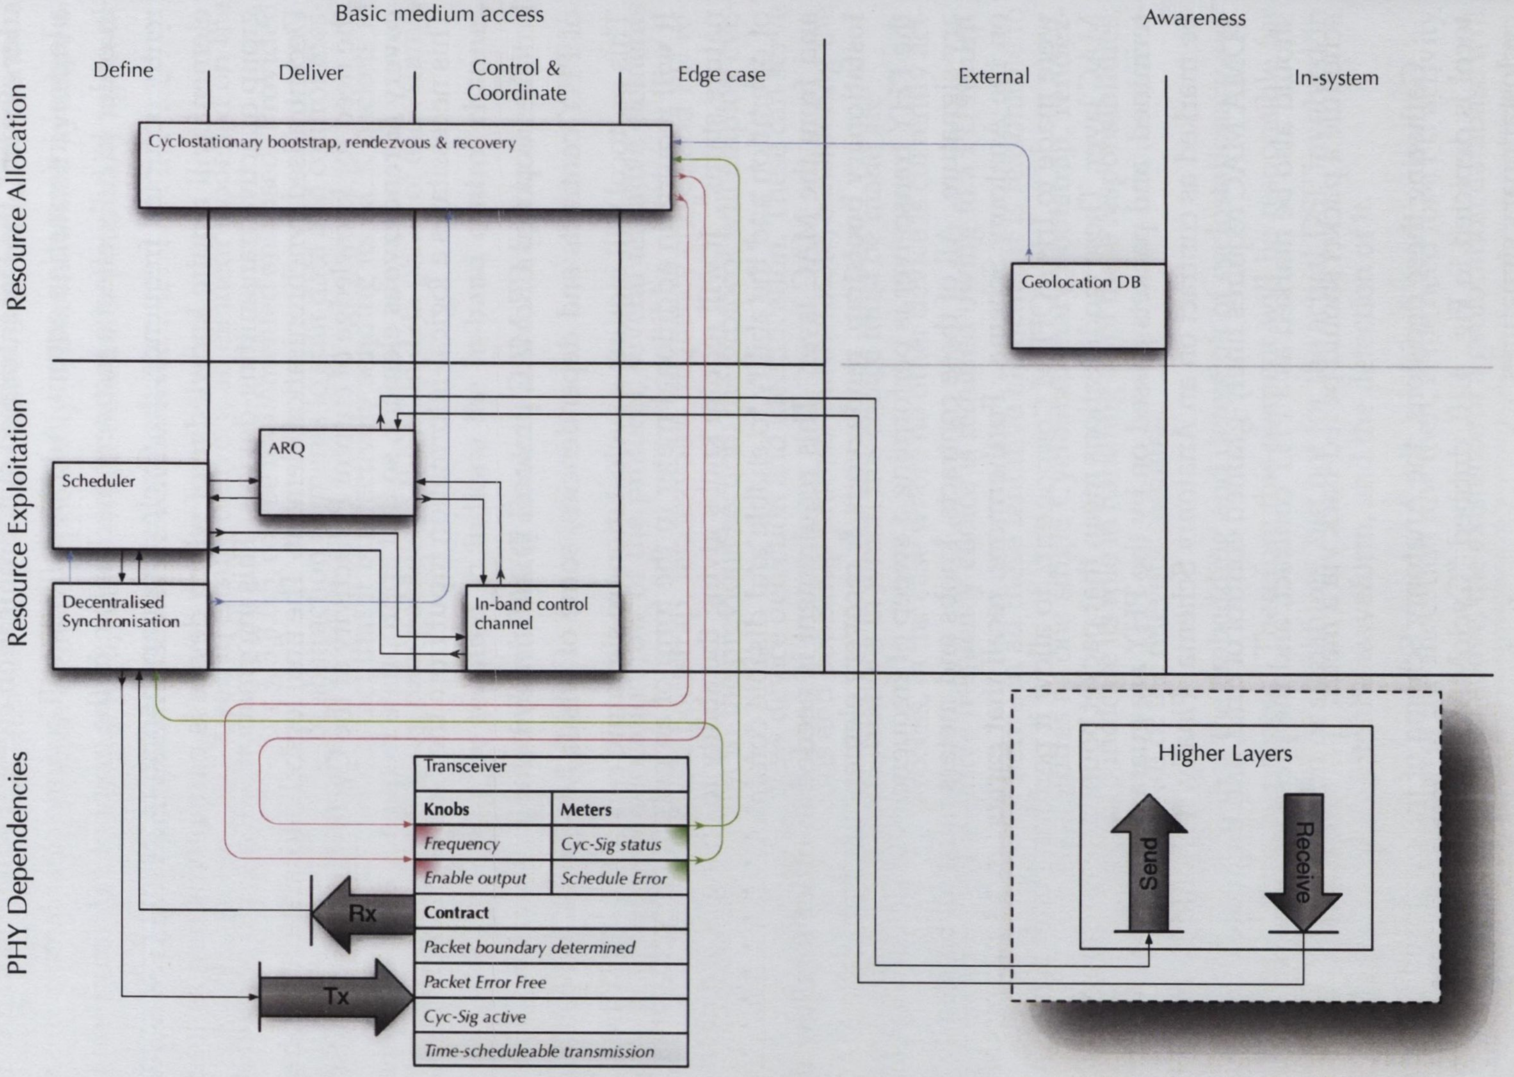

Figure 4.10: Anatomy Schema of the proposed CR MAC protocol.

4.2.4.2 Note necessary abstraction elements

In order to maintain separation between elements, **Control channel** and **Data channel** abstraction elements will be used to separate the Scheduler from the other MAC elements. This removes their dependency on this specific implementation of the **Scheduler**, simply requiring they be connected to any 2-pathway radio or MAC configuration. Fig. 4.11 shows the Anatomy Schema for the CycloMAC protocol redrawn to note the position of abstraction elements.

4.2.4.3 Map Anatomy elements to Iris modules

Inspecting the Schema created in Step 1, it is clear that there are likely to be several Iris `Components` & `Controllers` to be created.

The **Scheduler** element determines when precisely in a frame a transmission should occur. It queues packets sent from the **ARQ**, awaiting a messages from the **Synchronisation** element as to when to transmit. As such messages will be logical-information – i.e. on the control-plane in Iris – the **Scheduler** is an example of an element that will need to be implemented as both an Iris `Component` and `Controller`.

The **Decentralised synchronisation** element will provide the **Scheduler** with a control-plane signal to indicate the start of every frame. Furthermore, it will monitor a PHY meter which indicates scheduling errors at the PHY transceiver. As a PHY meter is logical-information, this element must also be implemented in Iris as a combination of `Controller` and `Component`. (See §4.1.3.3.)

As mentioned earlier, a simple **ARQ** element can be implemented, as a single Iris `Component`.

A simple **In-band control channel** element will be required to pass control data from the **ARQ** element to the **Scheduler** element. For this implementation it only needs to pass all received communications data through.

At the allocation level, a **Cyclostationary bootstrap, rendezvous & recovery** element must be implemented to act on cyclostationary signature information from the PHY, and make changes to the frequency of operation. This element depends only on logical-information. Its inputs are logical-information from the **Geolocation DB** element, and the cyclostationary signature status meter from the PHY transceiver. Its output is a selected frequency, and whether to enable or disable passing of communications data from the transceiver to the MAC. Because this element needs access to the control-plane in Iris, but not the data-plane, it will be implemented in Iris solely as a `Controller`.

Finally, the **Geolocation DB** element as marked on the Anatomy Schema also deals only with logical-information. As a result, it too will be implemented as a `Controller`.

4.2.4.4 Build Iris MAC

As noted in §4.2.4.3, there are several new elements to be implemented in Iris.

The **Scheduler** `Component/ Controller` pair must be implemented. It will queue received data packets, awaiting a `Command` which is triggered by the **Synchronisation** element. The `Command` contains the timestamp to be used for the next upcoming slot. As it requires `Commands` – which operate on the control-plane in Iris – from a synchronisation element to indicate the start of a frame, it must be implemented in Iris as a combination of a `Controller` and `Component`.

A **Decentralised synchronisation** element must also be implemented as a `Component` and `Controller` pair, as it has control-plane input from the PHY. To implement it, a cluster-based decentralised approach was chosen, in which one device is elected to provide a common reference to all others. Should this device cease to function, another may take control.

Each device is given a unique node ID. To commence, each node sends synchronisation request packets at random intervals, while listening in the intervening period for the sync packets of others. If after a period of time no synchronisation packet has been received with a higher node ID than this device's ID, it deems itself elected *cluster-head*. The cluster-head begins to broadcast beacon packets to all neighbouring devices, indicating the start of a *slot*, which is a group of frames. Each device can calculate its slot within this frame as a function of this ID. Any device that fails to receive a beacon will restart this procedure after a timeout, and should a cluster-head hear beacon packets of a higher ID than its own, it shall cease transmitting beacons and calculate its time slot relative to the device with the higher ID. (See Fig. 4.12. In this figure, the cluster-head is referred to as *MASTER* and the all other devices as *SLAVE*.)

The **ARQ** element is implemented as a `Component` using the simple description of ARQ from Chapter 3, 3.3.0.5. The abstraction components being used mean that it will be potentially reusable in many future protocols.

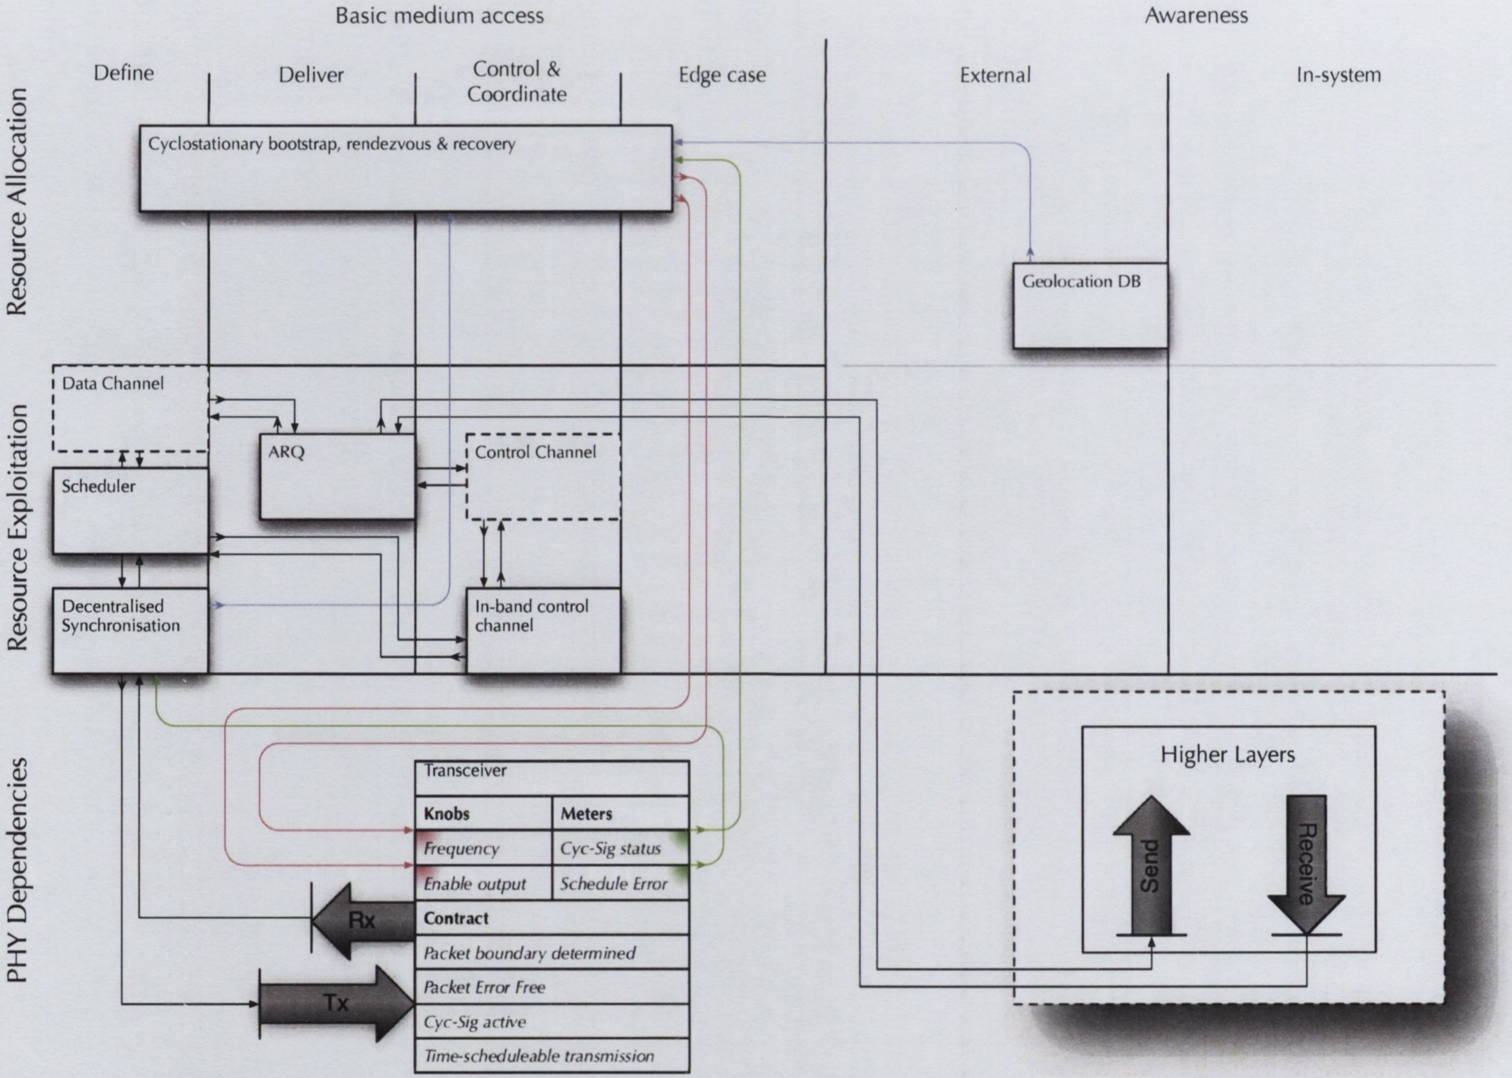

Figure 4.11: Anatomy Schema of the proposed CycloMAC protocol, noting abstraction elements.

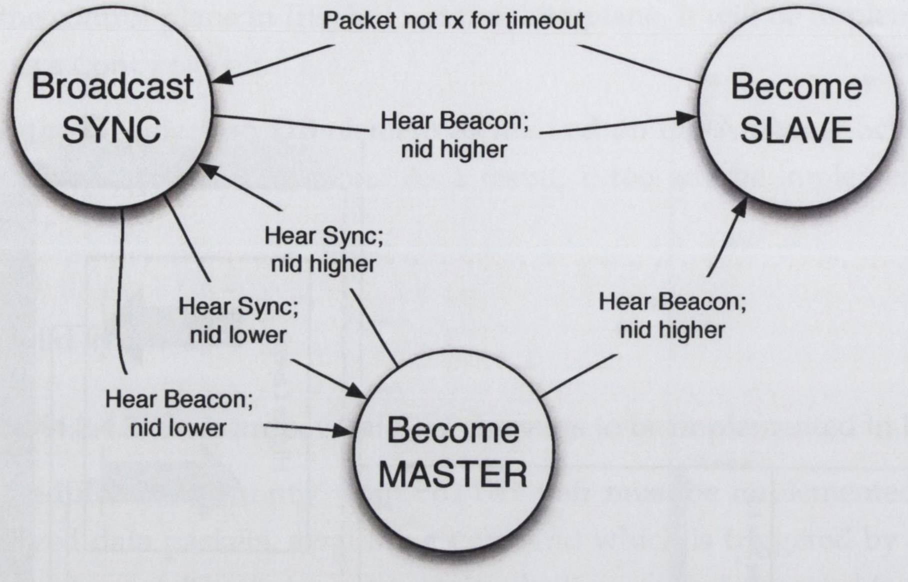

Figure 4.12: State machine of the cluster-based **Decentralised synchronisation** element that enables distributed TDMA in the CycloMAC protocol. In this figure, the cluster-head is referred to as *MASTER* and the all other devices as *SLAVE*.

An **In-band control channel** element's `Component` must also be implemented. As it needs to do no more than pass through data from the **ARQ** element's control output to the **Scheduler** element's control input and vice-versa, its implementation is trivial.

Next, the **Cyclostationary bootstrap, rendezvous & recovery** element's corresponding `Controller` must be implemented. Its responsibilities are as follows: i.) When activated, it must determine if any other devices are already active; ii.) If none are, it must select a frequency to operate on and ensure devices activated subsequent to now can detect it; iii.) If it does detect another device or devices, it must tune to their frequency to allow other elements to begin to communicate and iv.) It must move and recover communications in the event of a PU arrival on the currently occupied frequency.

The element was implemented to behave as follows: When activated, communications data-output from the transceiver is disabled, and the cyclostationary element instructs the PHY to iteratively sweep across the entire range of possible frequencies, to check for the presence of a signature of pre-agreed cyclic-frequency. If no signature is detected after a random but bounded period of time, the PU activity logical information input from the **Geolocation DB** is checked, and a random centre-frequency not currently overlapping with a PU is chosen. It instructs the PHY to tune to that frequency and enables passing received output to other elements. If a signature is detected, the element will request that the PHY tune to that centre-frequency, and enable passing received output to other elements, at which point they will begin normal operation. In the context of the **Decentralised synchronisation** element implemented above, this will mean the device will start transmitting intermittent synchronisation request packets. Owing to the *Cyc-sig active* PHY contract, a cyclostationary signature will automatically be embedded in these transmissions.

A subsequently activated device will perform the same sweep across frequencies, and detect this signature. It will then tune to the appropriate centre frequency, and begin to interact with the initially activated device by electing a cluster-head, as per the procedure outlined above.

In the event of a PU arrival, the element will be made aware of it by its constant monitoring of the input from the **Geolocation DB** element. In this case, communication must be ceased, immediately. This implementation takes the strict view that even a single transmission intended to aid in the communications recovery process that knowingly interferes with PU transmissions is not permissible. Given this, the choice of what frequency to move to must be made in some way that does not cause fragmentation of the communicating group, and that can be shared in order to resume communication. The simplest method allowing this is for one device in the group to "lead" the move. If the **Cyclostationary bootstrap, rendezvous & recovery** is aware by means of logical-information input from the **Decentralised synchronisation** element whether it is a cluster-head or not in the context of the TDMA, this already-made decision can be employed. The elected cluster-head can therefore make a frequency selection in the same manner as for the bootstrap process above[9], assured that no other device in the group will also attempt to make a decision. It tunes to this chosen centre frequency, and recommences the synchronisation process. All non-cluster-head devices cease transmission, and recommence sweeping for the cyclostationary signature of the cluster-head's signal. Once detected, communication can be resumed.

The final element to be implemented in Iris is the **Geolocation DB**, which will be implemented as a `Controller`. As noted earlier in §4.2.3.5, the precise method by which a Geolocation database is communicated with by a CR, and how CR obtains its location are beyond the scope of this thesis's implementation. Therefore for this

---

[9] *i.e* Choose a frequency at random with recourse to current PU activity as indicated by the **Geolocation DB** element.

implementation the implemented `Controller` reads from a text-file which represents current PU activity in the spectrum, and may be edited to represent a change in same.

With all elements now specified and implemented in Iris, the CycloMAC can now be built.

4.2.4.5 Build Iris PHY

The next step is to create a compatible PHY layer. It can be seen from the PHY contract in Fig. 4.10 that there are six additional PHY contracts and dependencies beyond those offered as part of the OFDM `Component` discussed earlier, in §4.1.2. These are: the ability to time schedule, presence of a cyclostationary signature, a meter to indicate a scheduling error, a meter to indicate detection a cyclostationary signature, the ability to control the transceiver's frequency and the ability to enable and disable output from the PHY, which prevents the MAC receiving spurious input when the signature detector is sweeping across frequencies. Therefore the same OFDM transceiver can be employed, once these differences have been addressed.

The OFDM modulator and demodulator discussed earlier in §4.1.2 was modified to incorporate cyclostationary signatures.

In order to detect the presence of a signature, a pass-through `CyclostationaryDetectorComponent` which already forms part of the Iris library, is inserted into its receive chain.

Existing `UsrpTxComponent` and `UsrpRxComponents` expose frequency as an Iris Parameter, and therefore can be issued `ParametricReconfigurations` by the suitably designed `Controller` part of the **Cyclostationary bootstrap, rendezvous & recovery** element.

In order to allow that same element's `Controller` to disable PHY output while sweeping through frequencies to detect signatures is in progress, it can take advantage of an existing parameter in the `CyclostationaryDetectorComponent`, which controls whether communications data is passed through or absorbed. This prevents spurious data from being passed to the demodulator and MAC while a signature remains undetected.

These represent most of the steps needed to allow construction of a PHY transceiver appropriate to this CR MAC. The remaining steps relate to time scheduling, and will be discussed in the subsection below.

A note on SDR-specific implementation challenges

As is well demonstrated in the literature, a primary challenge in implementing TDMA protocols using discrete radio front-ends connected to standard GPP-based PCs is the unpredictability of the connecting bus, in addition to the non-realtime nature of the GPP itself [75–78]. When these two factors combine, there can be a significant variation between when a packet is expected to be modulated onto the air and when it actually occurs, which is not predictable packet-to-packet, and even less so device-to-device.

An alternative is to take advantage of specific newer-generation USRP functionality which allows blocks of complex base-band samples to be sent to the USRP with a timestamp indicating when they should be transmitted. These samples are queued in a FIFO buffer until the appropriate time – or should that time already have passed by the time the block reaches the USRP, they are dropped.

As part of the distributed TDMA proposed above, the elected master device sends beacon packets at frame-length intervals, determining the interval by calls to a microsecond accuracy local timer, using the `microsec_clock::local_time()` function from the boost C++ library's `posix_time` functionality [100]. As there remains the possibility of the host-clock falling behind that of the USRP, the USRP will pass an error using its asynchronous transmit-error queue, back to the host. This may be used by any scheduling code to correct its time error. In the context of Iris, such an error can used to generate an `Event.` A `Controller` can be used to respond to such an `Event.` In this case, a `Controller`, considered a part of the **Synchronisation** element, allows it to act on and correct these errors in scheduling ahead.

This scheduling feature of the USRP hardware can therefore meet the required PHY contract . Furthermore, the USRP output of scheduling errors acts as the required PHY meter linked to the **Decentralised Synchronisation** element shown in the CycloMAC Anatomy Schema.

4.2.4.6 Create Iris xml

Fig. 4.13 represents the implementation in Iris of the proposed CycloMAC protocol. The xml version is simply a linearised representation of the same connections and can be seen in Appendix A.1.

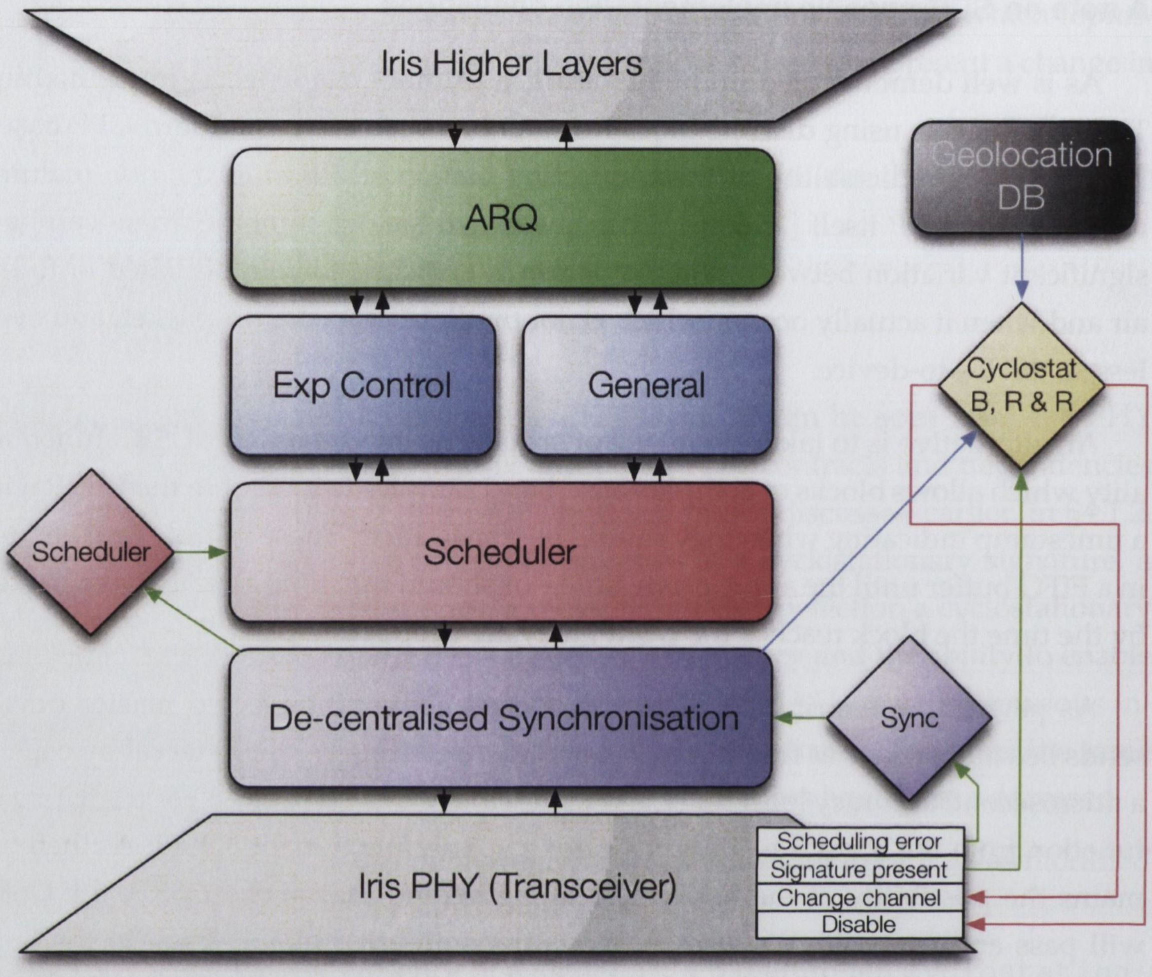

Figure 4.13: The Iris representation of the proposed CR MAC protocol.

4.2.4.7 Run

With the xml complete, it can be passed to the Iris core for testing and experimentation, which will be discussed in the next section.

4.2.5 A note on real-world implementation and experimentation

It should be noted that in addition to the significant technical hurdles to be overcome in performing any real-world radio experimentation, and the reasonable amount of expertise required, a particular problem with real-world experiments is reproducibility and meaningful quantitative results. On the one hand, experimentation is one of the fundamental principles of science, and with engineering defined as the application of science to build systems, real-world implementation of a system is the implicit end-goal. On the other, the nature of the hardware and systems involved in experimentation for CR research is such that the deduction of meaningful and transferable quantitative results is difficult. As SDR platforms run on GPPs on a non-real-time OS, using real-world manufactured and hence non-identical hardware, in constantly varying physical and RF environments, every single experiment on every device is on some level unique. Nychis *et al.*exemplified this in their measurement of the inconsistent and unpredictable behaviour of USRP hardware [77].

As discussed earlier in Chapter 2, §2.4.1, the majority of published research on MAC protocols – both for CR and for legacy devices – makes use of simulation as its evaluation tool. While simulation offers the potential for reproducibility and the ability to verify basic performance, there are significant failings in simulation as is actually performed in the literature, where aids to reproducibility, tables of parameters, and in some cases even details of the platform employed are unavailable. This was comprehensively documented for Mobile Ad-Hoc NETwork (MANET) research by Kurkowski *et al.*, but the same shortcomings apply in CR and CR MAC research [32]. This is of course notwithstanding the fact that it is impossible to simulate every variable offered in the real world, and to simulate real hardware, so even simple real-world experimentation can offer insight that a very advanced simulation cannot. Therefore while the majority of the literature's evident preference for simulation is understandable, this thesis would argue that even if the current failings in much of its implementation were remedied, the importance of real-world experimentation as proof-of-concept would not be discounted.

4.2.6 Experimental behaviour and results

Using USRP N210 hardware, transmitting at 5.003GHz using XCVR2450 daughterboards, the implemented CycloMAC was tested using the parameters listed in the table in Fig. 4.14b. A photograph of the CTVR testbed is shown in Fig. 4.14a.

To examine the over-the-air transmission behaviour of the system, results and plots were recorded using a Rohde & Schwarz FSVR Spectrum Analyser [101].

In order to examine the signal structure for correct TDMA behaviour, Fig. 4.15 plots the time-domain power from 3 devices using CycloMAC. The TDMA frame structure can be observed clearly from this figure. The elected master's short beacon pulse indicates the start of a frame, and it is followed by a control subframe. Following the control subframe, the data subframe contains a data packet transmitted by each device, spaced very closely in time, but not colliding. Note that it is the USRP's

(a) Photo of CTVR testbed, featuring 4 Iris-GPP-USRP nodes.

Center Frequency	5.003 GHz
Bandwidth	1 Mhz
Modulation Scheme	OFDM
Active data subcarriers	192 of 256
Subcarrier Modulation	BPSK
ARQ timeout	50 ms
Data Packet size	734 B

(b) Table of RF parameters used.

Figure 4.14: Photo of testbed setup and table of RF parameters employed.

time-scheduling feature that allows for this very fine-grained and consistent time-domain performance, which is accurate to ±1 sample. For this experiment's sample rate of 1MS/s, this results accuracy of $\pm 1\mu$s.

It should be noted that all devices consistently have packets to send. Were this not the case, the slots assigned to some devices might go idle, even as other devices with large queues of packets could potentially make use of them. More advanced MAC protocols may allow more flexible access to the medium, to take the nature of device traffic into account in making best use of the medium.

Overall, the system was found to perform as expected. Multiple devices initialised were capable of bootstrapping to the same frequency, electing a tempor-

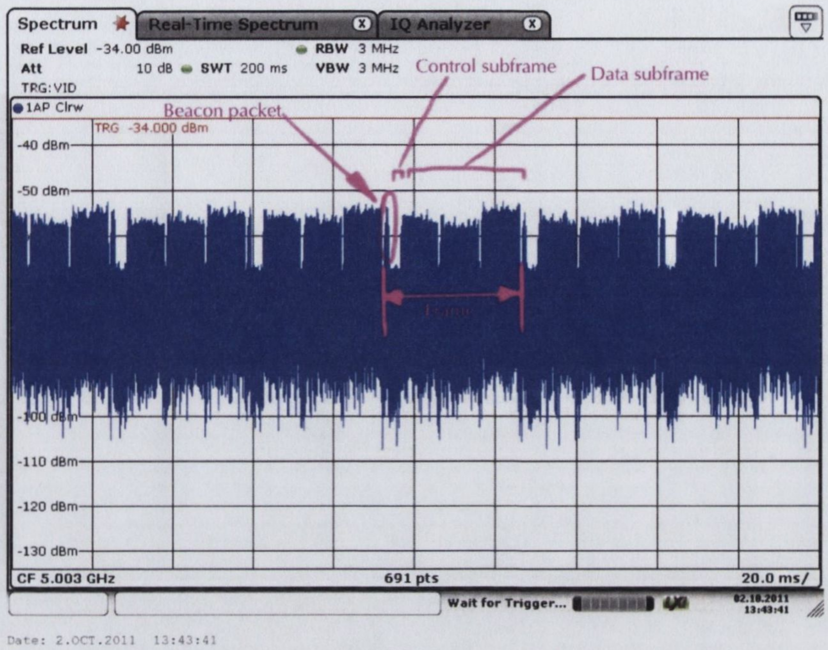

Ref Level -34.00 dBm
Att 10 dB SWT 200 ms
RBW 3 MHz
VBW 3 MHz
TRG: VID
1AP Clrw
TRG -34.000 dBm
CF 5.003 GHz
691 pts
20.0 ms/
Wait for Trigger...
02.10.2011 13:43:41
Date: 2.OCT.2011 13:43:41

Figure 4.15: Time-domain amplitude plot of transmissions from 3 nodes employing the CycloMAC protocol, recorded from over-the-air transmissions. A complete TDMA frame is highlighted. Its control and data subframes are also marked. Note that for this capture, control subframes are empty as packets were transmit broadcast, and therefore do not receive an ACK.

ary cluster-head, and beginning transmission using TDMA. Subsequent devices of higher node id could then be activated, discover the networks, and following a brief handover period, be themselves elected cluster-head. Those same devices could be deactivated, after which the network would revert to using another device as cluster-head.

Finally, to consider a real-world instance of one of the interactions introduced in §4.2.2, the process of recovery from PU arrival on an overlapping frequency is examined: Following the commencement of transmission by a PU on a frequency overlapping with the CR MAC protocol, the system appropriately backs-off and resumes transmission on an alternative frequency. This process was captured using the Rohde & Schwartz FVSR Spectrum Analyser, with the results plotted in Matlab for ease of representation, and shown in Fig. 4.16. For this example the CycloMAC system is operating centred at 5.003GHz, while a PU of narrower bandwidth but higher power is operating at the adjacent frequency of 5.0015GHz. Following a transition of the PU system to centre on 5.003GHz, the CR system responds to the change by moving to an alternative frequency, at which communication is re-established. Note

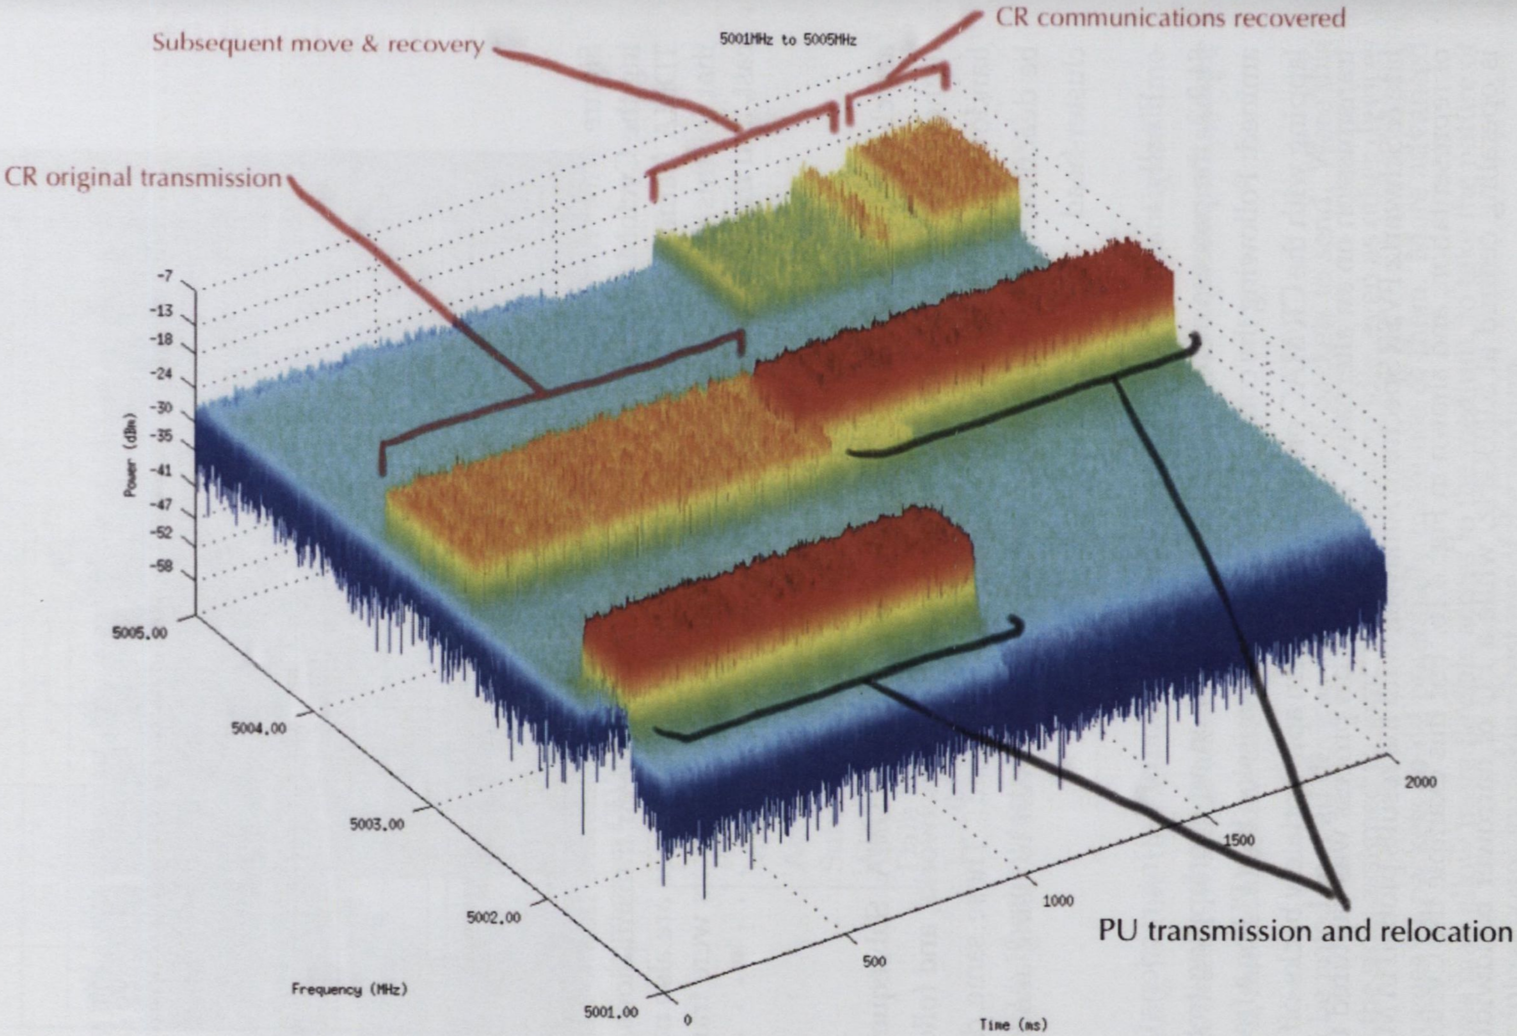

Figure 4.16: Over the air spectrum capture of the PU arrival recovery process for CycloMAC. As annotated in the figure, the CR MAC protocol changes frequency and recovers communication, briefly after the PU device's commencement of transmission on the same frequency.

that this process takes a moment, and corresponds to the time taken to reelect a temporary-master and re-establish full TDMA communication.

4.3 Summary

This chapter demonstrated the Anatomy of a CR MAC in action. It first introduced the Iris platform as a candidate for an example implementation of MAC protocols using the Anatomy. It then explained a step-by-step method for implementing a CR MAC protocol represented using the Anatomy in Iris. It demonstrated this method by employing it to create an advanced CR MAC protocol – CycloMAC – to suit a DSA scenario. It then detailed the implementation of this protocol in Iris and examined the performance of the system, while noting the difficulties in drawing results from experimental work. The following chapter will discuss this implementation in detail, examining the effect of using the Anatomy to implement this CR MAC, making observations on future options for expandability this affords, and explaining how the Anatomy addresses the limitations with current methods of CR MAC design, as motivated in Chapter 2.

5 Discussing the Anatomy in action

The previous chapter introduced the design and implementation of Cognitive Radio (CR) Medium Access Control (MAC) protocols using the Anatomy, doing so on an example platform. This chapter will discuss the Anatomy of a CR MAC in broader terms. It will discuss the use cases of the Anatomy in detail, with further examples to reinforce the previous chapters. It will consider how it addresses the issues raised in Chapter 2, and the concerns raised in Chapter 3. And It will then explore its relationship with other flexible, cognitive and extensible MAC research. This work was introduced in Chapter 2, §2.4.3 and will here be revised in the context of the Anatomy.

5.1 Use cases of the Anatomy of a CR MAC

This section will discuss the uses of the Anatomy of a CR MAC under the two broad headings of deconstruction and reconstruction, representing as they do the core concepts it aims to facilitate.

5.1.1 The Anatomy for deconstruction

The word *deconstruction* is here used in the broadest possible sense, to mean understanding, implementation, comparison and dissection of CR MAC protocols.

Consider first the great aid the Anatomy can offer to the understanding of individual protocols, and the comparison of multiple protocols. Chapter 2, §2.5 explained that the fundamental medium access options available to a wireless device are in fact not enormously varied or complicated: The medium may be divided in time, frequency or code. Simultaneous transmissions at overlapping frequencies by multiple devices in close proximity must be prevented or controlled. And extra information can be transmitted in addition to the core data being communicated in order to aid with addressing, data–integrity protection, and generally ease in running. This small range of options can be arranged and combined in an enormous number of ways, so the end result can certainly be complex. With MAC for CR, the complexity has the potential to blossom exponentially.

The Anatomy is first and foremost a schematic representation of a CR MAC protocol designed to capture and assist in managing this complexity, and provide a consistent mechanism for considering functionality in terms of the definition of a MAC introduced in this thesis. Recall this defines a MAC as responsible for:

1. Choosing when, where and how to access the spectrum resource/ medium.
2. Ensuring message delivery.
3. Handling edge-cases of the chosen access mechanism.
4. Arrange necessary control signalling.

The Anatomy recognises that a MAC can be considered a combination of *elements* rather than a monolithic block of functionality, and therefore mandates a deconstruction of a protocol into these constituents. The constituents are the smallest reasonable standalone parts of a MAC that potentially offer original functionality. This rough level of granularity being advocated – though not mandated – reduces the complexity involved in understanding a protocol, while still accurately representing it.

There is a large degree of disagreement in terms of the accepted definition of a CR MAC, varying from an extremely flexible protocol, to a slightly modified legacy one. This is in addition to a consistent conflation of CR with opportunistic overlay Dynamic Spectrum Access (DSA). The Anatomy provides a rapid and accurate mechanism to see what the proposer of a protocol actually means. The distinction between a *resource allocation* and a *resource exploitation* protocol is a particularly important one to make. The Anatomy clearly allows for this, while noting most importantly that the two need not be – and indeed should not be – mutually exclusive. Though examples of both variety of protocol are numerous in the literature, emphasis and awareness of this separation by authors is not. Given a significant proportion of proposed CR MAC protocols are resource allocation protocols only, and hence may not be considered MAC protocols at all by some traditional measures[1], recognition of both of these roles is an important contribution of the Anatomy. As a corollary to this point, it should be noted that in some cases, authors of resource

---

[1] For example, if one were to define a MAC protocol as "the full specification of transmitter and receiver actions required to effect communication", a resource allocation protocol would fail to meet the definition.

allocation protocols do not themselves consider them to be CR MAC protocols. As a result they may not be labelled as such at all, and may be referred to in the literature as channel selection mechanisms, dynamic spectrum allocation mechanisms, or other titles. Once employing the definitions called for in this thesis however, all such protocols fall under the umbrella of resource allocation in the Anatomy Schema, recognising their contribution as an important but not necessarily lone portion of a CR MAC.

The current dominance of simulation as an evaluation tool is a matter discussed in Chapter 2, §2.5. This thesis argues that implementation and experimentation – aided by simulation for extrapolation of large scale performance – should be the ultimate test and goal of any CR MAC research. The Anatomy aids in implementation of proposed protocols by mandating consideration of the physical (PHY)-layer requirements of any protocols. These take the form of its *dependencies* – the knobs and meters with which it violates conventional cross-layer boundaries – and its *contract*. Therefore it aids not just in high-level understanding, but in understanding to a degree that should aid in real-world deployment of a CR MAC protocol.

This requirement for a detailed consideration of the MAC relationship with the PHY will also alleviate some of the issues associated with simulation, by increasing awareness of what must actually be simulated. Simplification of PHY behaviour will be of significantly greater difficulty to justify with an Anatomy characterised MAC. Furthermore, this accurate characterisation will allow simulations comparing MAC protocols to be as meaningful as possible, and avoid the pitfall of misunderstood and/ or conflicting assumptions by authors.

Indeed, one of the greatest strengths of the Anatomy of a CR MAC is in its role as a consistent framework for comparison. When deconstructed to the element level, common features between CR MACs become much more readily apparent. As usefully, the degree to which a comparison is meaningful also becomes much more readily apparent. That is to say, MACs with largely similar features, but relying on very different PHY-layer dependencies and contracts could very well be more different than expected. This is helpful in determining avenues for reasonable future research, and in indicating its difficulty.

In a world where an Anatomy Schema accompanies every MAC design, capability is easily judged, radio requirements are always considered, and functionality and complexity are rendered clearly. Even when CR MAC is conflated with overlay opportunistic DSA, the Anatomy makes precise functionality clear, so the limitations of inaccurate labelling can be side-stepped.

Of course it is not just useful to look at a deconstructed CR MAC as a whole of parts. The elements can also be examined individually. The importance of this feature cannot be overemphasised, as it provides a method to extract useful features from what at times can be monolithic protocols. These monolithic protocols may be of little novelty overall but contain small individual features of significant usefulness. In a world where elements are considered reasonable contributions alone, research can focus on the creation of new individual CR MAC elements, without needing to worry about the extraneous remainder of the MAC, knowing that others will be capable of incorporating this new element into their designs, as long as they are conscious of the interconnections of the element with other elements, its PHY dependencies and PHY contract requirements.

Even should work not be situated in the most general and useful manner, the procedure to create an Anatomy Schema from Chapter 3, §3.3 can be followed. Then those examining the protocol at a later date can extract conceptual MAC elements as desired, getting straight to the parts that are of genuine interest as quickly as possible. Common design features might include the ability to borrow a CR MAC protocol's control channel, or its unique channel allocation model, while ignoring the remainder of the protocol.

5.1.1.1 Example applications of the Anatomy for deconstruction

To demonstrate some of the aforementioned uses of the Anatomy in deconstruction, this section considers the completed Anatomy Schema for two CR MACs selected from the literature. They have been chosen from the range of protocols reviewed in Chapter 2, §2.4 as they are protocols of appreciable novelty and with significant contrasts. The Anatomy Schemata discussed in this section were obtained by following the instructions as laid out in Chapter 3, §3.3.

Recall the "DCR-MAC" proposed by Yoo *et al.* was used as an example for creation of an Anatomy Schema earlier, in §3.3.0.8 of Chapter 3 [36]. The completed Anatomy Schema representation that was derived is reproduced here in Figure 5.1.

The Anatomy makes several aspects of this protocol clear at a glance. Firstly, it is a resource exploitation protocol. It therefore can stand alone to facilitate packet-to-packet communication between multiple devices without any further extension. Its lack of a resource allocation part however does mean that it may not be able to allocate the medium reasonably between larger groups of devices, and facilitate the division into subgroups. In such a subgroup, each device would be contending with fewer neighbours. Hence it would be likely to gain more frequent access to the spectrum resource, and potentially effect more efficient communication. Secondly, it has a default requirement of two transceivers, one with adjustable frequency, the other with a contract to be time-synchronised with its peers and permit the time-schedulable transmission of short radio frequency (RF)-energy pulses. Thirdly, its uniqueness stems primarily from its **Multi-stage channel selection handshake** mechanism, rendered as one element in this Anatomy Schema.

These three simple observations would not be so easily made without the Anatomy of a CR MAC. Nonetheless, between them they represent a significant insight into the functionality and capability of the protocol, what might be required to implement it, and what is novel about it.

To contrast with the above example, the "local bargaining" mechanism for DSA proposed by Cao *et al.* is also considered [51]. This protocol was discussed as part of the literature overview, in Chapter 2, §2.4. It is chosen here for discussion as it is a well cited but simple example of a resource allocation protocol. It offers an elegant representation of the network, representing potential interference between nodes on a conflict-graph, and provides an intuitive and distributed means by which nodes may heuristically bargain to ensure that no node on the graph is "starved" of access to channels.

More specifically, when and where to access the spectrum resource is determined by assuming it has been divided into fixed frequency channels, and with each node being assigned a set of channels. This is performed using a heuristic approximating optimal graph-colouring. Nodes control and coordinate via a Common Control Channel (CCC), starting with an arbitrary channel assignment. Using dedicated second transceivers, they perform "local bargaining" for channels with their immediate neighbours via the CCC, on an as-needed basis. A "poverty line" metric devised to guarantee a minimum number of channels per node is used to determine whether bargaining should take place, and whether a node might altruistically grant a channel to another node. Nodes periodically broadcast their spectrum assignment to neighbours in order to allow poverty line calculations to take place, and should any node be calculated as below it, it will commence the bargaining process. Ensuring delivery is not dealt with in as part of the protocol, and the edge cases of hidden and exposed nodes are assumed non-present by requiring the CCC transceiver have a transmission range of twice the data transceiver. It should be noted that once assignment has taken place, the details of when and how access is performed on a message-to-message basis is not defined, making this very clearly a

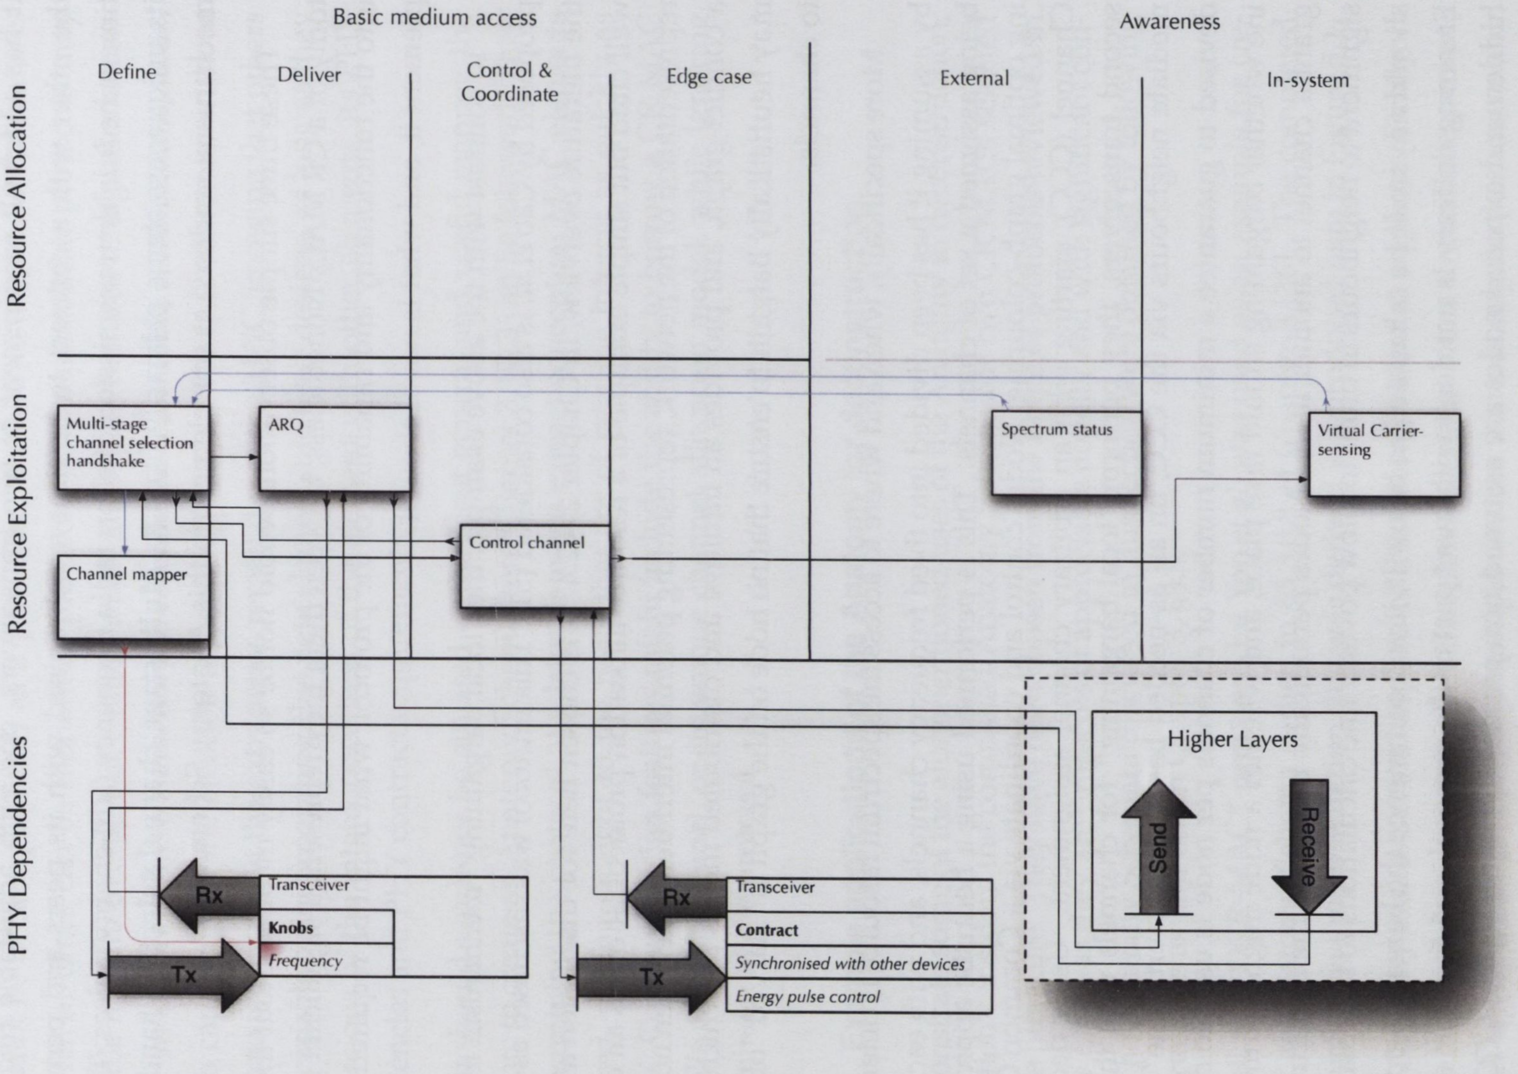

Figure 5.1: Completed Anatomy Schema for "DCR-MAC" proposed by Yoo *et al.* in [36].

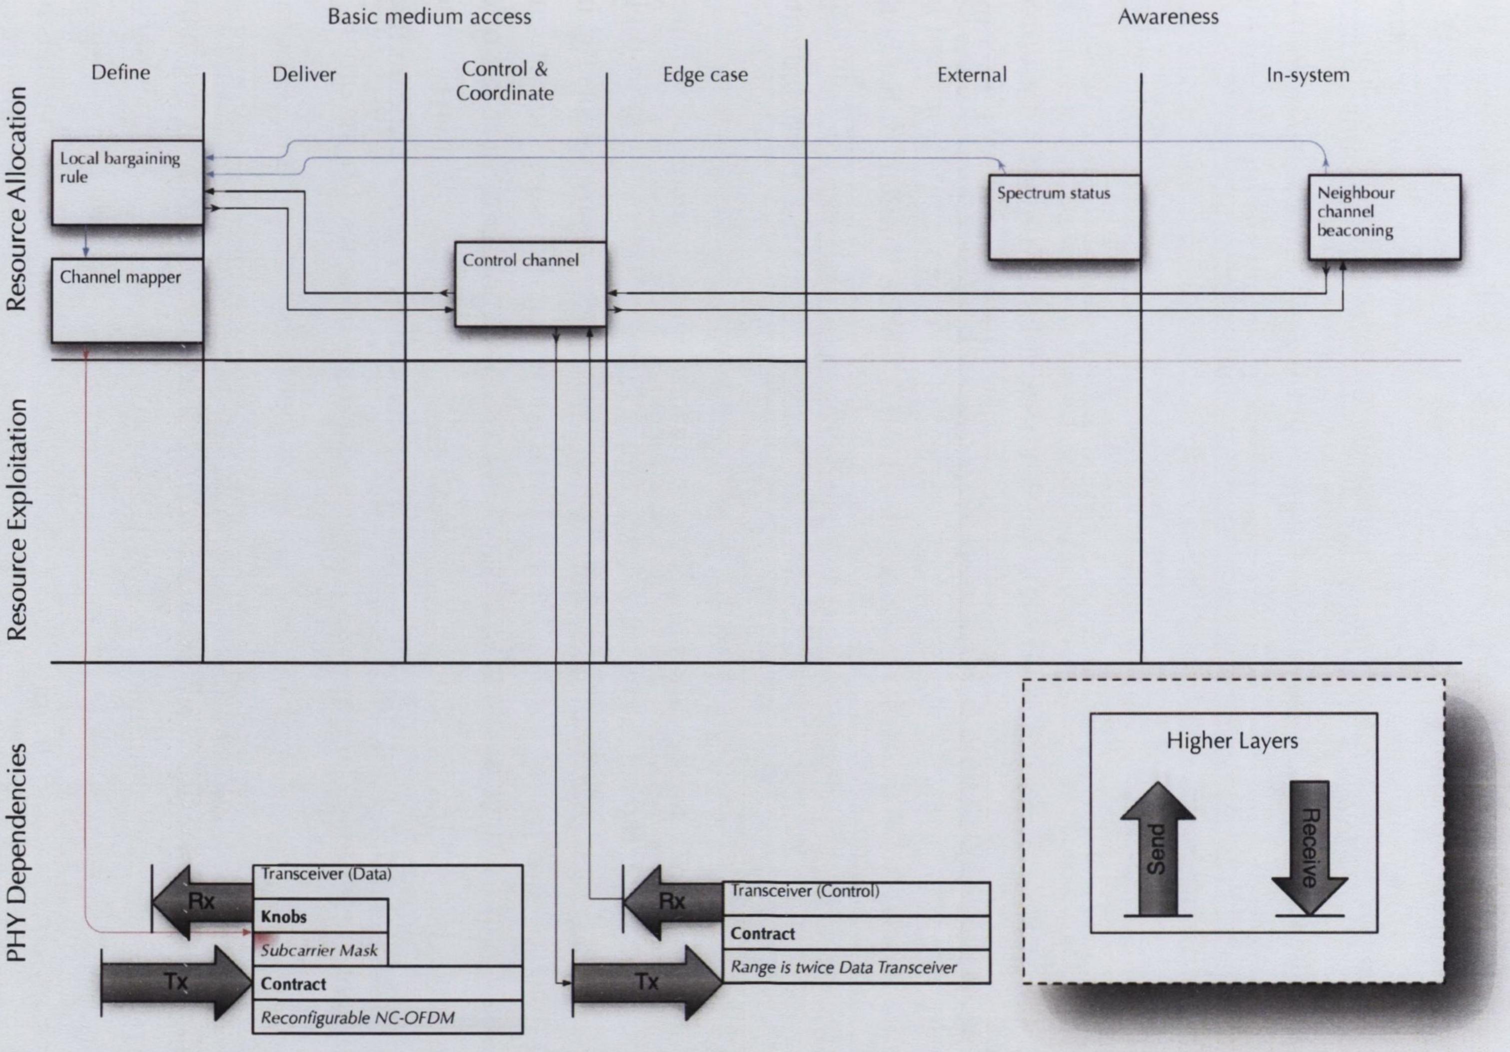

Figure 5.2: Completed Anatomy Schema for "Local bargaining MAC" proposed by Cao *et al.* in [51].

spectrum resource allocation protocol.

A completed Anatomy Schema corresponding to it is shown in Figure 5.2.

This Anatomy Schema again allows for several observations at a glance. The first is that there is no pathway for data provided from higher layers to the PHY – this protocol is purely focused on resource allocation. The second is that it centres on the control of a PHY capable of non-contiguous wideband transmissions[2], which clearly requires some sort of resource exploitation mechanism. All possible bandwidth coverable by this transceiver represents the resource being allocated. And the third is that it requires a dedicated CCC.

These two short examples have shown how the Anatomy can serve as an aid to understanding, but with both to hand, comparison of them can also be discussed.

The differences between the two protocols are immediately obvious: resource allocation and resource exploitation are – as explained in Chapter 3, §3.1.2 – fundamentally different responsibilities, and therefore this serves to limit the extent to which the protocols in their entirety can be compared. They are similar in that both require two discrete transceivers, albeit with substantially differing PHY contracts and knob and meter dependencies. They further share in the fact that both dedicate one for use on a CCC.

Examining on an Anatomy Element level, a further common feature becomes clear. Both require information as to the activity of Primary Users (PUs) in their neighbourhood – their **Spectrum status** elements. Both protocols are unspecific in terms of their specification of such an information source, and therefore the element representing it on their Anatomy Schemata serves as a placeholder for a real-world implementation which may be more complicated, and might require its own tie-in to transceiver(s) and other Anatomy Elements. The Anatomy Schemata indicate that such a source implemented for one of these protocols may be of use for the other.

It is also worth pointing out another pair of elements that the protocols appear to share, their respective **Channel mapper** elements. In each protocol the purpose of this element is to take "channel" as an input from some decision making mechanism, and interpret it for manipulation of a PHY knob. However, though they serve the same effective purpose, their details are different, as the protocols do not agree on what they define as a channel. The DCR-MAC of Yoo *et al.* defines channels as frequency blocks, which the CR can tune to and use one of at a time. The local bar-

---

[2] Specifically, the authors intend use of a Non-Contiguous Orthogonal Frequency Division Multiplexing (NC-OFDM) PHY transceiver.

gaining **Channel mapper** of Cao *et al.* defines channels as the subcarrier blocks of an NC-OFDM PHY, and considers a CR device capable of using many not-necessarily contiguous channels simultaneously. Indeed, the purpose of making a **Channel mapper** an explicit element, rather than an assumed feature of all MAC protocols, is to acknowledge this potential divergence of definition.

Fig. 5.3 shows the Anatomy Schema for the DCR-MAC, but with its novelty highlighted [36]. As mentioned earlier, the protocol's primary novelty comes from its **Multi-stage channel selection handshake** element. This element receives data from the layers above, has a connection to a control channel and depends system awareness in the form of both current PU activity, and the activity of other Secondary User (SU) devices from virtual carrier-sensing. It must also have the ability to request a change of frequency of its PHY transceiver.

Fig. 5.4 highlights the novel elements in the Anatomy Schema of the Local Bargaining protocol [51]. Its **Local bargaining rule** element is the primary contribution of the protocol, but also distinguishing it is the in-system awareness mechanism. This involves each device's constant beaconing of detected channel availability and its own occupancy, and is marked as the **Neighbour channel beaconing** element.

The importance of the Anatomy of a CR MAC here is three-fold. It is able to represent the element(s) which self-contain the major novelty of these particular protocols. It also clearly shows their role in the overall CR MAC protocol. And it clearly indicates what other elements in a protocol they depend on.

Therefore, we can see that on many levels, the Anatomy can assist in the process of deconstruction of CR MACs.

5.1.2 The Anatomy for reconstruction

Reconstruction is also meant in a broad sense. It encompasses the thorough and forward thinking design facilitated by the Anatomy; the capacity to extend existing protocols by the addition of new elements; the combination of groups of elements representing new merged protocols; and the wholesale creation of new protocols from elements sourced from a variety of others, or newly devised as required.

When a designer has a specific and nuanced understanding of the scenario for which a MAC is designed, their choices may be so nuanced and specific as to render the protocol incomparable with other proposals. Though it cannot prevent scenario specificity, the Anatomy forces consistent representation. This serves both to assist

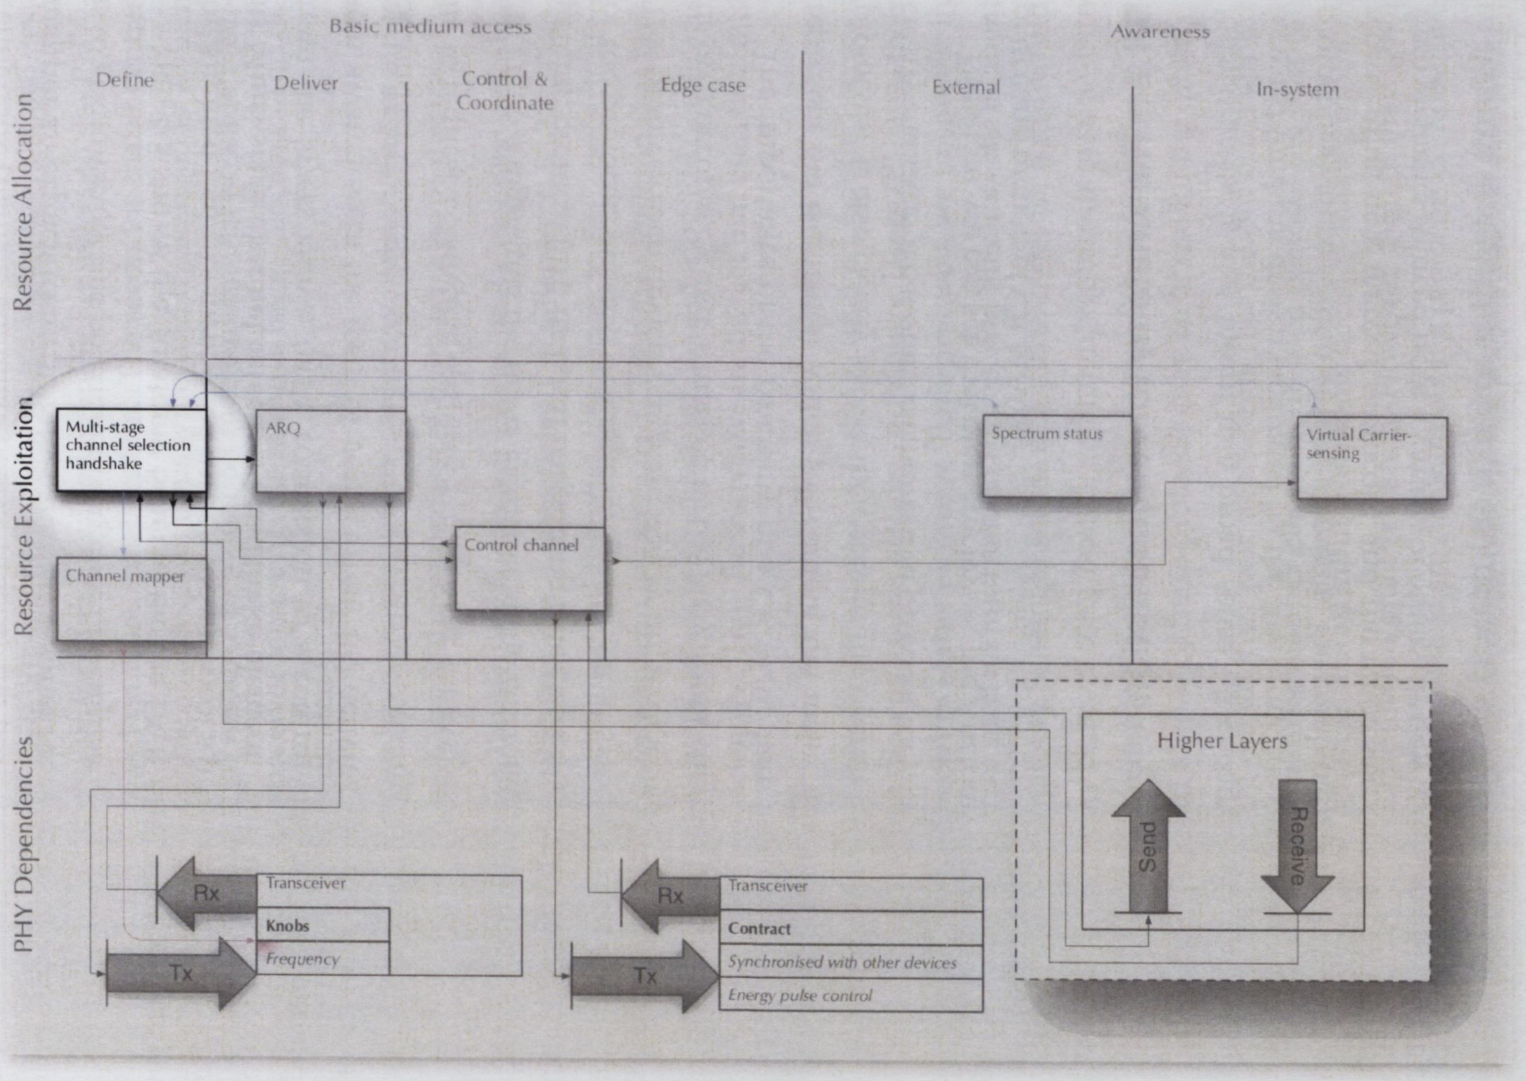

Figure 5.3: Anatomy Schema highlighting novel element of the "DCR-MAC" proposed by Yoo *et al.* in [36].

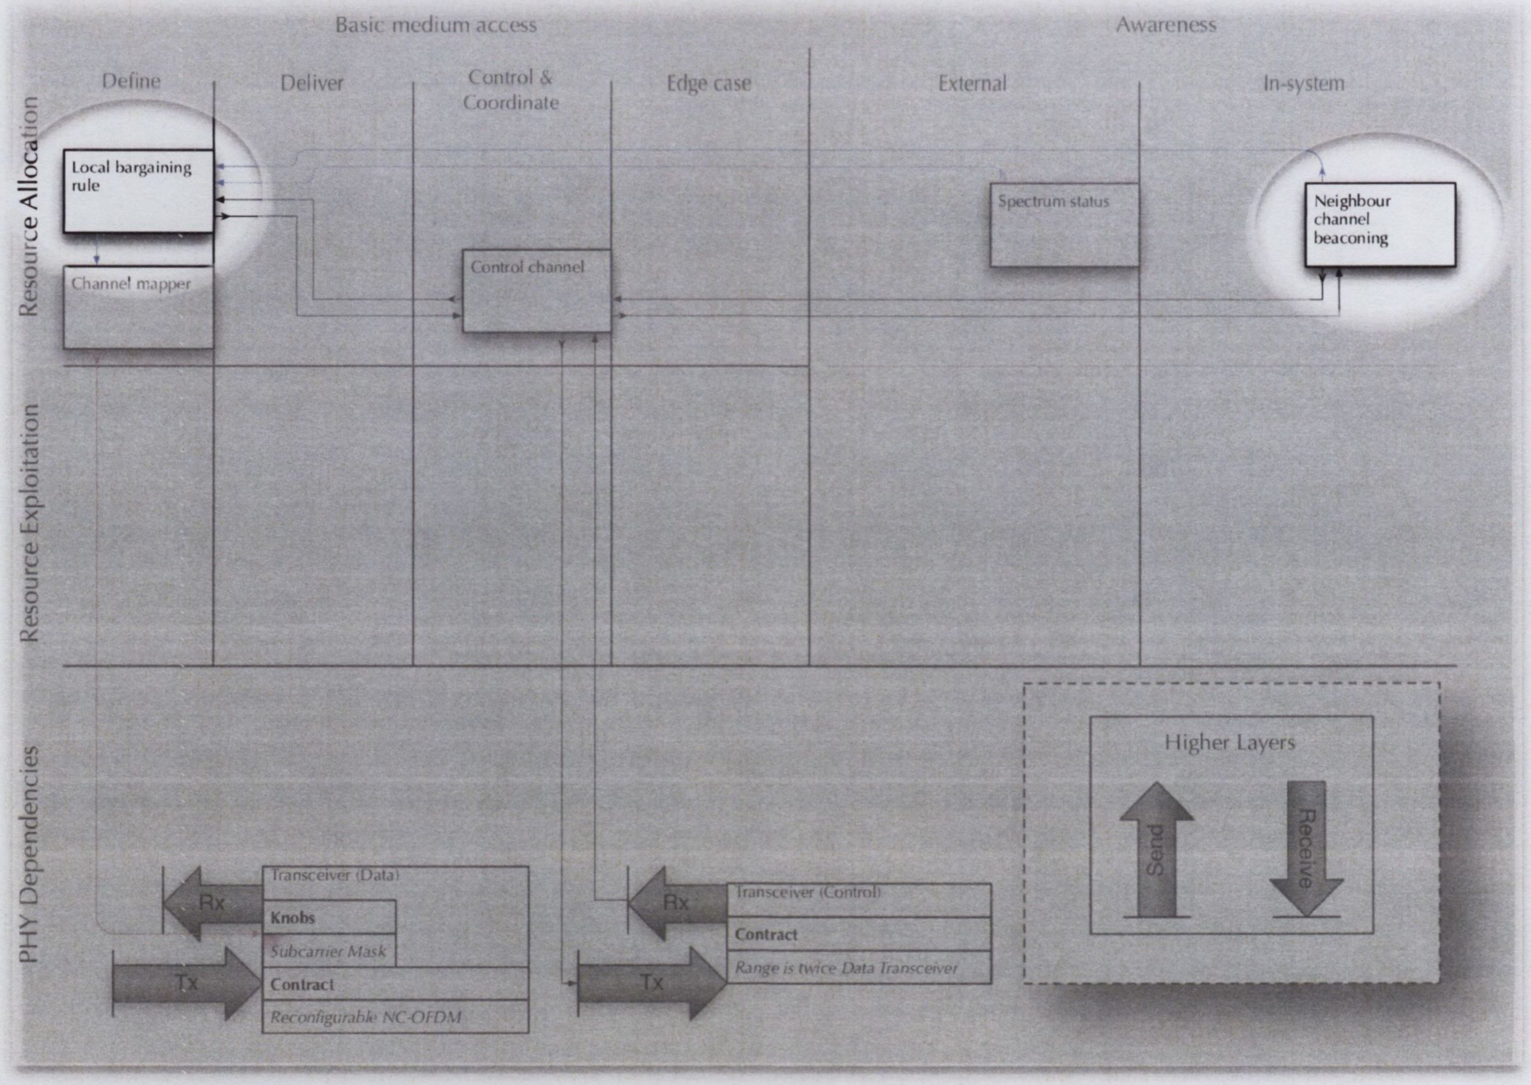

Figure 5.4: Anatomy Schema highlighting novel features of the "Local bargaining MAC" proposed by Cao *et al.* in [51].

others in the manner explained in §5.1.1 and to serve as a constant reminder to the designer of their assumptions, which might easily be forgotten. More profoundly, forcing a designer to represent a protocol in a manner consistently over time may reveal unexpected relationships and commonalities between protocols, and foster innovations in terms of both design efficiency and awareness of possible redundancies.

This can be seen if we consider some canonical simple MAC protocols: Aloha, a Round Robin (polling) protocol and a simple centralised TDMA protocol. The differences in the functionality of the overall protocols bely their structural similarities, which can be seen when they are laid out as a progression as in Fig. 5.5. When implemented using the Anatomy, each requires only a relatively small and simple series of modifications to become its more complex neighbour. In this way we can trace the "implementation evolution" of leading from a protocol as simple as Aloha, to an advanced protocol like CycloMAC. Aloha is composed primarily of an **Automatic Repeat reQuest (ARQ)** element, and optional abstraction elements. Replace its **Mux** element with a **Poller** which polls devices for packets however, and it becomes a Round Robin protocol. Replace that **Poller** element with a **Scheduler** and **Synchronisation** element then, and it becomes a Time Division Multiple Access (TDMA) protocol. From there, only limited modifications and additions are required to produce CycloMAC. The **ARQ** element can remain unchanged, even as the elements around it increase in complexity. Furthermore, the **Scheduler** element can be shared between CycloMAC's more complex distributed TDMA, and the simpler centralised TDMA protocol. In short the Anatomy both helps clarify existing work, and point out meaningful paths for future work.

There is a strong visual appeal and intuitive element to the somewhat organic approach to CR MAC design permitted by the Anatomy – a MAC's functionality may grow, being added to element by element. Use of abstraction elements to make elements agnostic of their surrounding neighbours means that conceptually large changes become easier, and facilitates this growth. Simply put, by making designers think about CR MAC elements modularly, they increase reuse and help reduce conceptual complexity. A control phase can be replaced with a control channel by replacing the connection to a scheduler with a link to an independent transceiver. Multi-channel communication can be multiplexed to single channel communication. And specific elements related to in-system or external awareness can be easily replaced or modified.

A number of international research groups are focused on the design of the net-

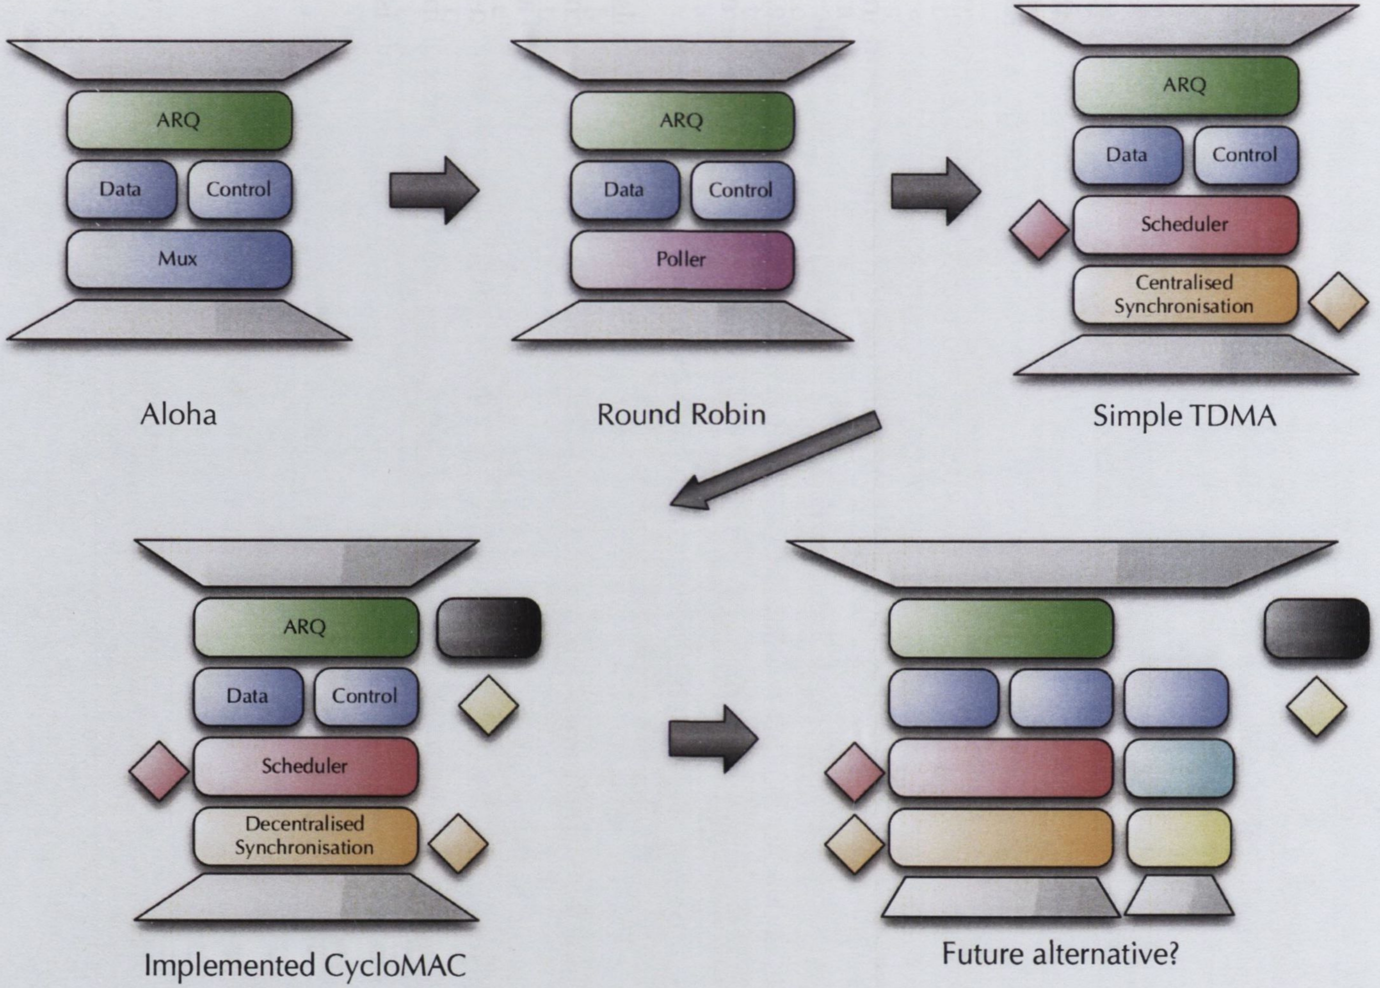

Figure 5.5: A simple figure illustrating the reusability of Anatomy MAC protocol elements, and of the ability to "collapse" one protocol configuration into a differing one with very simple changes to its Anatomy configuration. The figure shows the Iris outline of 3 simple protocols, progressing in order of complexity. It then progresses to the implemented CR MAC, and a fifth non-specific protocol, to demonstrate how further significant design changes still allow for element reuse.

works of the future, including my own, CTVR [102]. One acknowledged method for doing this is to build networks that are evolvable. Taking the bigger picture view of network infrastructure from core to edge, fixed and wireless, it is impossible to design a system that is provably – or even likely to be – precisely sufficient for future needs. The best way to attempt to do so is by creating as evolvable an infrastructure as possible, leaving specification as unfixed as possible in order to allow future engineering to adapt the infrastructure to as yet unforeseeable needs, while requiring minimal or no changes to existing users and applications. The idea of "architectural anchors" is that flexible design, open to future innovation, is best performed by building around those portions of infrastructure most difficult to change, but making no other assumptions [103]. This occurs at several levels: A core that is resilient to changes in expected data traffic and volumes, in the manner that fibre networks have been thus far, must be designed. There must be a move to a metro-to-edge network that steers clear of expecting specific behaviour, but instead facilitates possible changes in traffic volume, source, patterns and network ownership. And in wireless, the Anatomy of a CR MAC brings this same thinking to CR MAC, by describing protocols and protocol elements in as "unanchored" a manner as possible, while still recognising that anchors – such as PHY dependencies and contracts – are very likely to be present for specific protocols and their elements.

In this unanchored environment, resulting elements of deconstructed protocols can be combined to form completely new CR MACs. A library of elements, each offering features and limitations, can be combined to tailor a CR MAC solution to a specific scenario with ease. All that is necessary is a clear awareness of its interconnections, dependencies and contracts, for an element to form part of a protocol in an unforeseen manner, logically very similar but functionally very different. This is among the most powerful features of the Anatomy of a CR MAC.

5.1.2.1 Example applications of the Anatomy for reconstruction

Chapter 4, §4.2 provided a full application of the Anatomy of a CR MAC for design of a protocol. In the case of the Anatomy Schemata of the DCR-MAC and Local Bargaining MAC discussed earlier in §5.1.1.1 (Figs. 5.1 and 5.2), the Anatomy was obviously not applied as part of the design process by the authors of the represented CR MAC protocols. Nonetheless, once Schemata have been created, they unquestionably facilitate forward-thinking implementation. Two issues alluded to earlier, the control channel requirement of both protocols, and their **Spectrum status** element are items that could benefit from a forward-thinking design. Design of elements that are not dependent on how the control data-paths they connect to are implemented would allow their future use with different mechanisms. And implemented in a suitably generic and self-contained manner, a common PU information gathering method could be used for both protocols, whether it be a completely new design, or extracted from another protocol.

Moving further to consider the idea of wholesale reconstruction – the combination of groups of elements from different protocols – the two Schemata presented earlier in the chapter are now re-examined. As the Local Bargaining Protocol Schema in Figure 5.2 is a resource allocation protocol, and the DCR-MAC, Figure 5.1, is a resource exploitation protocol, the two might be combined largely intact to produce a new protocol capable of both.

The process of combination of protocols is a visually intuitive one, but it should be noted that it does present a risk, in that it may at times give an illusion of excess simplicity. This is a problem caused by the sharing of terminology between roles in a CR MAC Anatomy Schema, and in particular can arise when an unwary designer does not take care with cross role interactions. The primary rule in order to avoid the risk of an inconsistent combination is to remember that all Anatomy elements should be considered by default to act only as part of *one* of a CR MAC's roles.

Figure 5.6 shows an Anatomy Schema of a protocol incorporating both the DCR-MAC of Yoo *et al.* and Local Bargaining protocol of Cao *et al.* Only minor modification is required to combine them both. They may share a common **Spectrum status** element as an information source. Following an allocation by the **Local bargaining rule** element of a group of NC-OFDM blocks, the frequency choice portion of the DCR-MAC's **Multi-stage channel selection handshake** element outputs a channel on which to communicate. In this case however, the channel will now represent the subcarrier block for an individual device's transmission. Put succinctly, the **Local bargaining rule** allocates the resources for the entire group of devices, and within the group (non-static) pairs that wish to intermittently communicate arrange it using the CCC and then do so using agreed portions of the allocated resource. As discussed earlier when comparing the two protocols in §5.1.1.1, both share a need for a dedicated CCC. By using a scheduler and a PHY contracted to provide time-schedulable transmission capability, both protocols may share a single PHY transceiver and channel for control, splitting one control channel into two logical ones.

The paragraph above noted that complexity in the combination of protocols relates primarily to cases where elements are shared between the two roles of a CR MAC. In this example, there are three such shared elements.

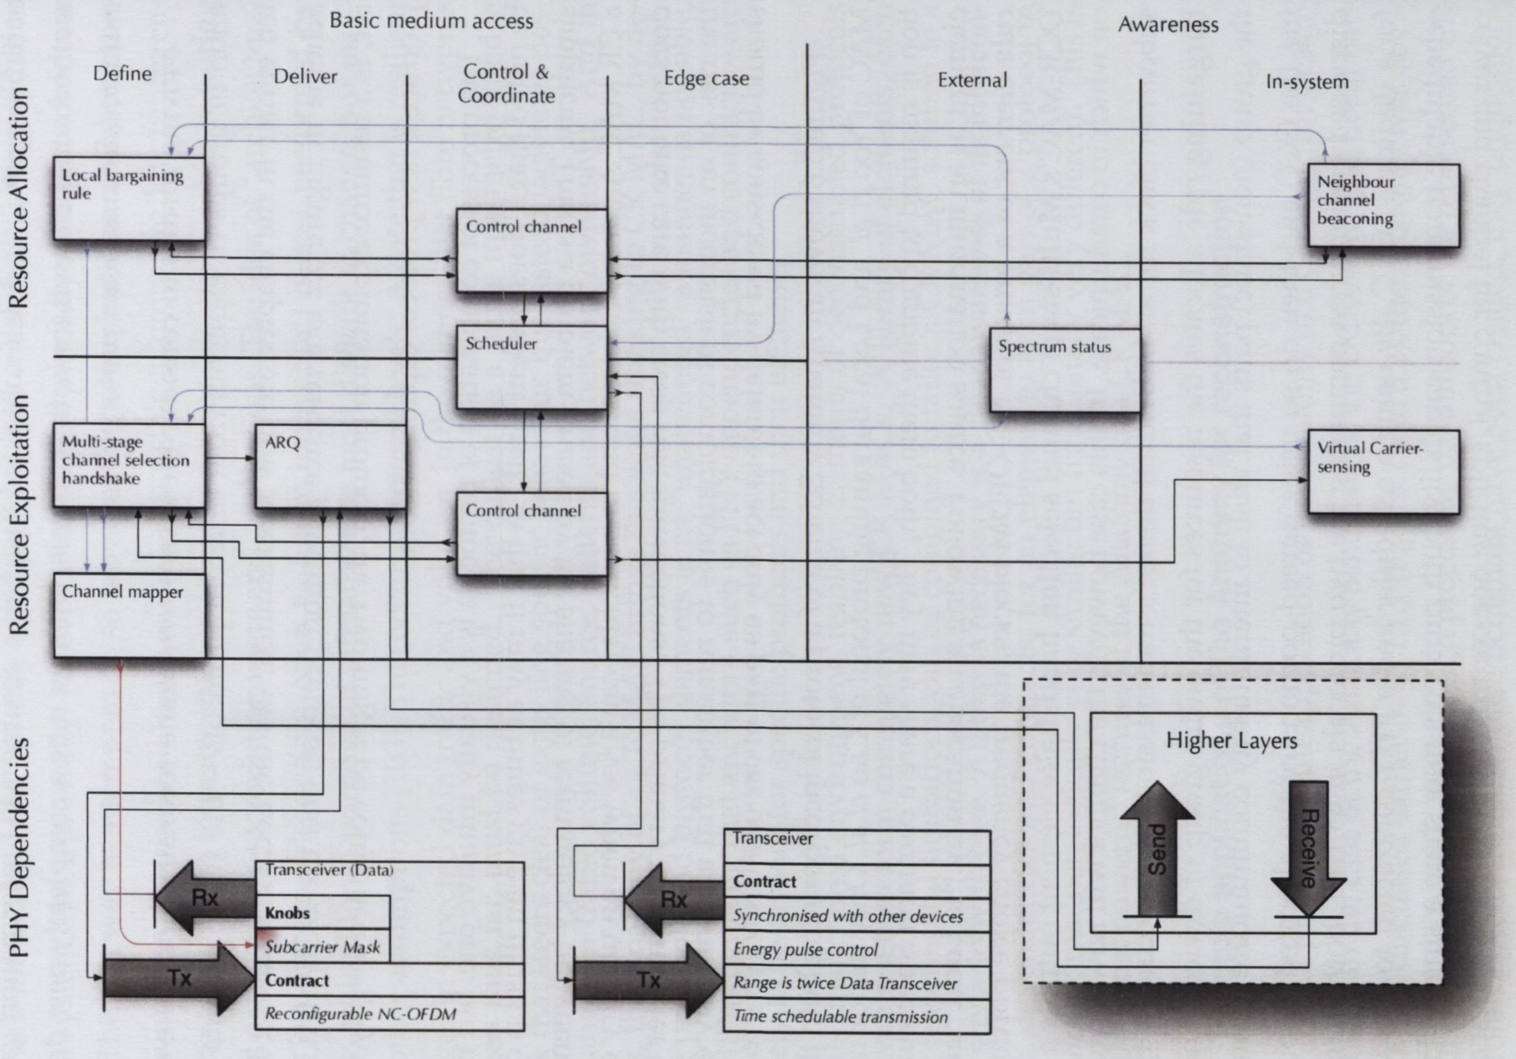

Figure 5.6: Anatomy Schema showing a combination of "Local bargaining MAC" [51] and DCR-MAC [36].

The first, the **Spectrum status** element, does not bring any challenges to the combination, as there is no scope for the elements part in one role to affect its part in the other. From the point of view of each role, its connection to the other may as well not exist – it is neither an absorbing nor a passthrough element.

The second, is the sharing of a **Channel mapper** element. In the context of this new combination protocol, the **Multi-stage channel selection handshake** element still outputs a choice of channel, which is passed to a **Channel mapper** element. The **Channel mapper** element however, is now shared between the allocation and exploitation roles of the protocol, and thus aware of the allocation decision made by the **Local bargaining rule** element. As a result, it can map the channel exploitation decision of the handshake element to the subset of subcarrier blocks allocated by the **Local bargaining rule**.

The third, is the sharing of a **Scheduler** element, in order to allow both roles to take advantage of the same CCC. It is important to stress that such a sharing of a transceiver is not ordinary behaviour. The responsibility for coordination and control as part of a MAC's allocation role is separate from that responsibility as part of its exploitation role. Therefore a control channel at one role is a conceptually separate entity from a control channel in the other. In fact, the set of devices being communicated with on each is very likely different. Thought of in terms of the example of an infrastructure network, the role of the allocation control channel is to allow communication between Access Points (APs), while the exploitation control channel is used by an AP for communication with its client devices. Every AP requires its own unique such control channel[3]. As a result, if the protocols of Cao *et al.* and Yoo *et al.* are combined but in a manner sharing a CCC as they are here, additional thought will be required to ensure that each role's control communication is not damaged.

The output of the **Local bargaining rule** element taken from the protocol of Cao *et al.* is a set of channels, passed to the **Channel mapper** to be translated into a corresponding subcarrier-mask for an NC-OFDM capable PHY transceiver. Once a mask has been allocated for a group of devices, their general data communications should occur on those subcarrier blocks only. In this new protocol, one transceiver is for data communications, and a second used jointly for resource exploitation control and resource allocation control communications. For resource allocation con-

---

[3] It may also be of assistance in this context to consider alternative but synonymous terminology for the allocation and exploitation roles. They can be considered as inter-network and intra-network communication, respectively, or as relating to co-existence and cooperation.

trol communications to occur successfully, all allocating devices within range of one another must be able to communicate with one another. By contrast, resource exploitation control communication should only happen within the group of allocated devices. There are two readily apparent solutions to this problem.

The first is to mandate that the scheduler further divide the control channel subframe dedicated to exploitation control communications, so that each group communicates in its own slot without risk of collision with out-of-scope control packets from other groups. The alternative is to employ a second NC-OFDM capable PHY for the control transceiver. Control communications during the resource allocation slot would be transmitted using all subcarrier blocks, and received by all devices in range. During the resource exploitation slot, each device would employ a mapped subcarrier mask corresponding to that allocated to its group for the data channel In this way, exploitation control messages are again restricted only to the group to which they apply.

For the protocol shown in Anatomy Schema form in Figure 5.6, the former choice is employed. The joint-control PHY transceiver is unmodified, but the scheduler will slot the exploitation control subframe. In order to do so, it will require knowledge of the number neighbouring groups, but this information is already available, as the **Neighbour channel beaconing** element provides this to the **Local bargaining rule** to inform its allocation decision. This means this requirement can be handled as a logical information input to the **Scheduler** element, as is marked on the Schema.

A second example is to consider a protocol combination involving the "CycloMAC" protocol designed in Chapter 4, §4.2. Its resource allocation portion consists of the **Cyclostationary bootstrap, rendezvous & recovery** mechanism, but does not facilitate the sort of inter-group frequency allocation offered in the Local Bargaining protocol of Cao *et al.*

In order to integrate the useful elements of the Local Bargaining protocol, an independent control channel for the resource allocation role is needed. The existing implementation of the CycloMAC protocol already provides for this in the form of a control-phase, therefore the local bargaining protocol could make use of it and not require an additional transceiver. The primary complexity here is similar to above, in that a single transceiver is shared for data, resource exploitation control and resource allocation control communications.

In the first combination example above combining the protocols of Yoo *et al.* and Cao *et al.* two potential solutions to this problem were discussed. For this example,

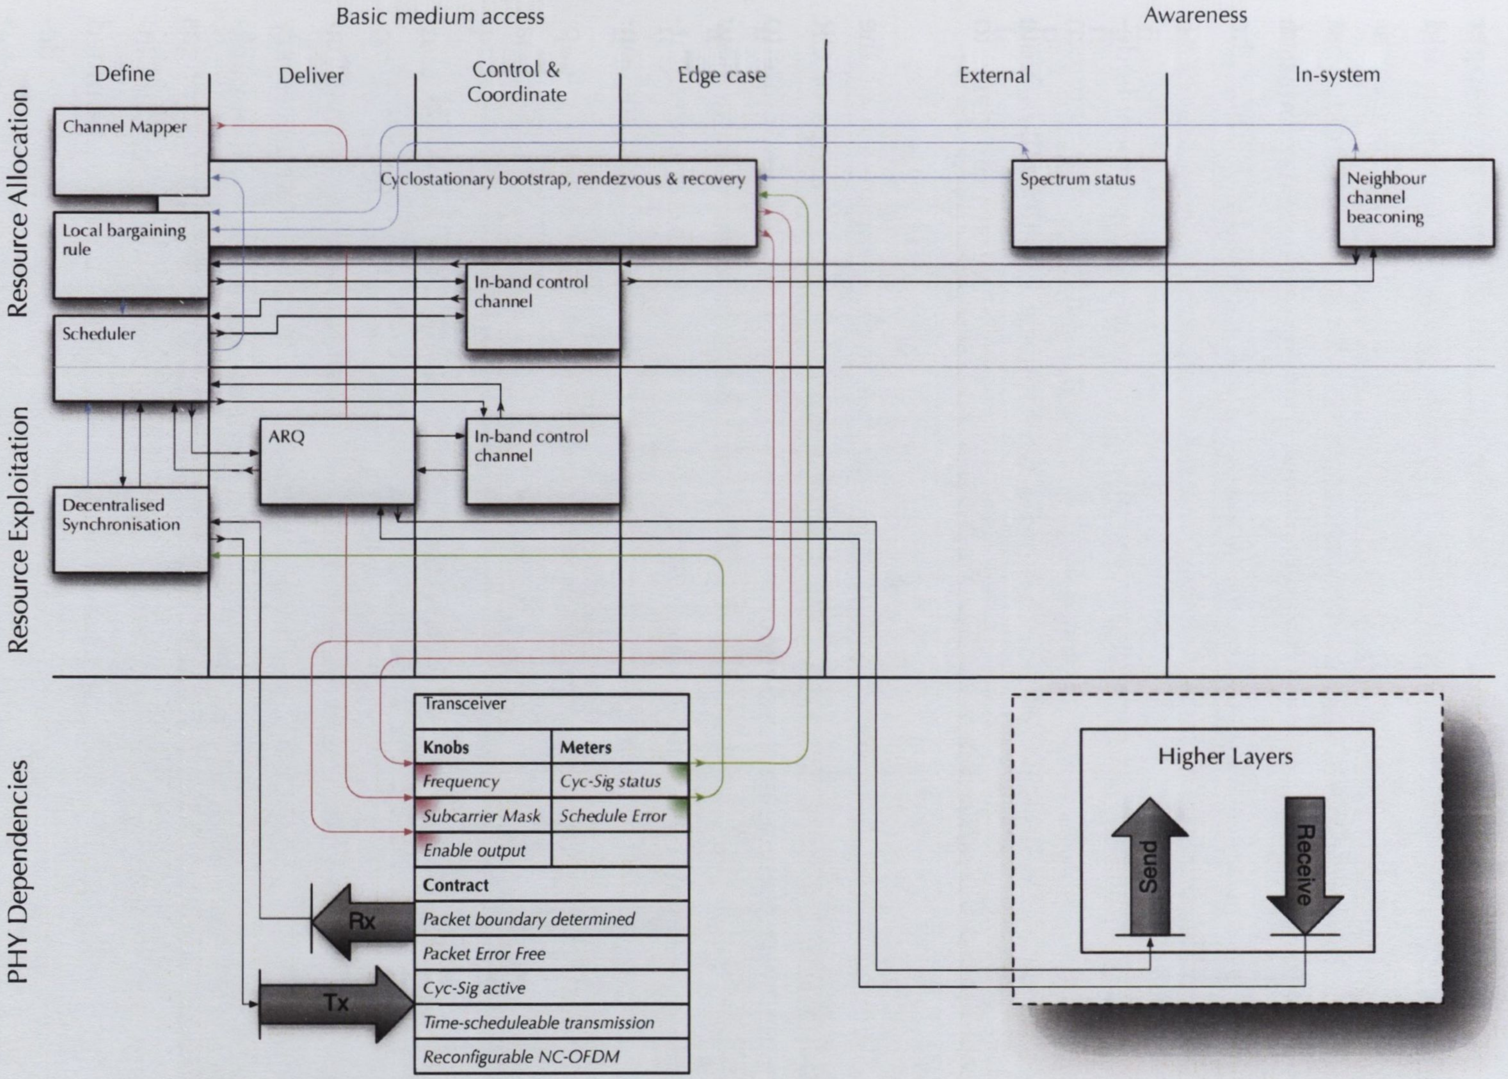

Figure 5.7: Anatomy Schema showing a combination of "Local bargaining MAC" [51] and the "CycloMAC" protocol designed in Chapter 4, §4.2.

the second of them will be chosen, namely that for the resource allocation control period alone, the allocating devices will need to communicate on *all* subcarriers, or an agreed subcarrier block. To implement these per-frame changes in the subcarrier mask, it makes sense to alter the Local bargaining protocol **Channel mapper** element's direct interface with the NC-OFDM PHY transceiver. Instead it should pass the subcarrier mask information to a modified scheduler, which can then itself interface with the PHY, setting the mask appropriately for each subframe. That is, the **Local bargaining rule**'s selected mask will be used for resource exploitation control and data subframes while the entire range of subcarriers can be used for the resource allocation control subframe.

The Anatomy Schema corresponding to the updated CycloMAC protocol is depicted in Figure 5.7. Note that space limitations necessitate the overlap between elements on the Anatomy schema, but this in no way reflects functionality. Note also the move of the Scheduler to cross the allocation/ exploitation boundary, to reflect that it contributes to both roles.

The combination of these two protocols results in a further expansion of the CycloMAC protocol to one with more nuanced behaviour, allowing groups of CR devices – perhaps representing independently owned networks – to cooperate to share the spectrum in a distributed and rule-based manner. It also shows how design using the Anatomy can help make future modifications like this one straightforward and easy-to-follow, even when they make significant changes to functionality.

5.2 Discussing PHY dependencies and contracts

Recall that Chapter 4, §4.1.2 noted that changes to the Orthogonal Frequency Division Multiplexing (OFDM) `Components` of the Iris software radio were necessary to allow for reasonable and efficient packet based communication i.e. so that packets of variable length could be transmitted. It further noted that additional changes were made, to perform Cyclic Redundancy Check (CRC) checks to ensure data integrity, and that the use of OFDM preambles was already automatically determined the start of a packet boundary. The end result is that the Iris OFDM `Components` determine packet boundaries via preamble synchronisation, determine precise packet length via included length header information, and confirm packet integrity via CRC checks, in a manner transparent to Iris `Components` above them.

When considering the example MAC protocols implemented in Chapter 4, these features of the OFDM transceivers are noted as PHY dependencies and contracts. Strictly speaking, these could be considered MAC layer responsibilities. Including them in a component here regarded as PHY represents a significantly blurred boundary in the stack. As discussed in Chapter 2, §2.2.7 however, blurring of these boundaries is extremely common. And though it may lead to potential issues in terms of CR MAC protocols silently developing dependencies on the presence of hardware, there are efficiency gains to be made by these sorts of integration. By making explicit a CR MAC protocol's contract with the PHY layer, so that its dependencies are made clear and considered when analysing the protocol, the benefits of such integration may be reaped without fear of a loss of awareness as to the full scope of what a protocol involves. This is especially necessary in cases where possibly unconventional action such as the extension or dissection of a protocol is taking place.

It may reasonably be asked what extra work would be involved for the protocols in Chapter 4 were they to employ a PHY incapable of meeting this contract. Anatomy elements responsible for CRC or a similar error checking mechanism would need to be created. And packet length information would need to be included in the protocol's frame format. Meanwhile, determining the start of a packet would require some sort of preamble mechanism to be present at the MAC layer. This known sequence would have to be correlated against all bits output by the now streaming PHY transceiver. Interestingly however, for packet-based communication, the corrections for frequency, timing and phase-offset recovery usually require a preamble-correlation process to be determined. If the MAC layer is performing the process, it must therefore feedback this recovery information in order to allow reasonable PHY demodulation performance over time. This means that the PHY is now dependent on the MAC to do its job, and shows that packet-boundary determination leads to unavoidable MAC-PHY overlap.

More generally, this shows that even where avoiding MAC-PHY overlap is desired – a decision that needs to be justified, as doing so may sacrifice efficiency gains – in wireless communication it is almost inevitable.

5.3 MAC-hardware dependency and Anatomy implementations

Recalling the discussion from Chapter 2, §2.2.7, once the separation between the MAC and PHY begins to blur, it becomes increasingly difficult to separate a protocol from specific hardware, or at least certain features of that hardware. A more general issue is that of the degree of disconnect between the PHY and MAC in hardware terms, most particularly in latency. Many MAC protocols, CR and otherwise assume a tight coupling and ability to very rapidly propagate commands to the PHY – with a similarly rapid turnaround for responses. This means that certain protocols are suited to system-on-chip type designs, or workarounds to approximate them.

PHY dependencies in the example MAC protocols constructed in this thesis have been restricted to scheduling in the case of TDMA, and to the reliability contract required by all implemented protocols. With the addition of a few more simple dependencies, significantly greater flexibility and advances in performance could be achieved for potential CR MACs. When remaining in the Software Defined Radio (SDR) domain, ultimately to move towards an ideal CR MAC, an at least partially FPGA assisted implementation of the Anatomy would be of use. In the TDMA implementation discussed in Chapter 4, §4.2.4.5, the scheduling feature supported by the Universal Software Radio Peripheral (USRP) 2/N2xx series hardware is an example of a partially FPGA assisted implementation.

This brings to the fore important questions as to the limitation of the Anatomy. It does assume a flexible, implicitly software-based design, allowing for reuse of modular elements. There is technically speaking nothing that prevents its use in a pure hardware implementation, but reconfigurability and reuse is a strong and important feature of it, which not all hardware is suited to taking full advantage of. A completely fixed implementation would represent a waste of possible functionality however, at least in the move towards an ideal CR MAC. Such a MAC is – as defined in Chapter 2, §2.3.2 one which can change its own configuration over time in response to environmental and internal stimuli. Partial reconfiguration in an FPGA would represent an avenue to do this, or at the very least run-time generation of the FPGA configuration, that the device may potentially improve its configuration iteratively, run-by-run. Owing to their flexibility, a well-designed FPGA implementation probably represents the best choice, perhaps building on a toolchain like TRUMP [62] (discussed in Chapter 2, §2.4 and later in §5.5).

5.4 Challenges faced by the Anatomy

Given the detailed discussion of the Anatomy of a CR MAC undertaken in this thesis, it is only right that it also be considered in terms of areas in which it is challenged. As a mechanism it is not without its limitations, and these will be explored in this section.

The first challenge to the Anatomy is the difficulty of constructing a Schema. While this process is obviously more straightforward when the Anatomy is used as part of the design process for a CR MAC protocol, accurately and completely characterising an existing CR MAC protocol can be an involved process, requiring a complete exploration of the protocol in question. This is an unavoidable difficulty.

The second challenge is the subjectivity of the choice of granularity. There is no specific scale of element mandated by the Anatomy of a CR MAC. This is a positive flexibility in one sense as it allows deconstruction to be performed at a level determined by the needs of the protocol designer or deployer. On the other hand, it leads to possible inconsistency of elements, and Anatomy Schemata representing the same protocols emerging radically different in surface appearance.

Extending this further, is the third challenge, that as a result of this subjectivity, there is no one "correct" Anatomy Schema representation of a CR MAC protocol. It should be pointed out that there are however accurate ones.

What might be considered a fourth challenge relates to the spectrum resource allocation role played by CR MACs, and is that of determining which device within each allocating group actually performs the allocation on behalf of the group. As discussed in Chapter 3, §3.4.2 however, this is not a limitation of the Anatomy, and is rather an anomaly of the mode of discussion used in the literature in proposing these protocols, which considers them as "nodes". This omission requires a CR MAC designer wishing to use the Anatomy to reconstruct a protocol playing both a resource allocation and exploitation role to define the means by which the "allocator" device is chosen. The extra elements required to facilitate such a combination were discussed and outlined in the aforementioned section.

A fifth challenge is the manner in which the design process of a CR MAC may become non-linear. In particular, when adding to an existing CR MAC protocol, but even possibly when reconstructing from scratch, the process may become an iterative, evolutionary process, with the satisfaction of each Anatomy element's requirements requiring a return to initial steps to add additional elements. Though this is not straightforward, and represents the altering of the design flows presented in Chapter 3, §3.3 and Chapter 4, §4.1.3 to become cyclical processes, this may be an unavoidable element of the process of design. It is also a realistic one, acknowledging as it does that feedback is can be a great contributor to design success, and that "single-pass" design is rarely capable of foreseeing all outcomes. It should be noted that iterative and evolutionary design processes dovetail with the greater motive of evolvable network design, as discussed earlier in this chapter in §5.1.2.

A sixth challenge relates to the points on subjectivity, but in harsher terms. Specifically, the Anatomy will not indicate if a CR MAC protocol is correct. It does not offer a means of verification of a protocol. This is a particularly strong challenge to the Anatomy, but in fact is a representation of the difficulty of verifying and predicting the behaviour of any complex system. It remains a requirement that a CR MAC designer take responsibility for design of protocols that perform as they wish to. The Anatomy assists by mandating that inputs, outputs and dependencies to and of elements be of the appropriate type, and all be satisfied, but ultimately, the designer must ensure that the resulting behaviour of the configuration is the desired one. It should be noted that this is a characteristic of modular design, in that it assists good design, but cannot prevent bad design. Furthermore, as noted by Partridge in his discussion of *"Forty Data Communications Research Questions"*, the formal verification of protocols is difficult problem. This is especially the case for automated verification, which is very much an open challenge in networking and communications research in general [104].

A related and seventh challenge is the characterisation of performance. One might well ask, given an Anatomy Schema, how to determine if it represents a "good" or "optimal" protocol. Or given a well-defined scenario, to determine which is the "best" or "optimal" CR MAC. Given a protocol solution, establishing whether it is anywhere on the scale from sufficient to optimal is a significant challenge. In addition to the challenges of verification noted in the sixth challenge, one or more objective functions are required, in order to provide some scale by which to define performance. Furthermore, performance with respect to the objective function must be measured or inferred. Measurement requires post-implementation assessment of real-world performance, or a method to infer performance of a design. Inference of the performance of an Anatomy designed protocol requires a priori knowledge of the individual performance characteristics of protocol elements, as well as well-defined rules for the combination of these characteristics in order to anticipate the overall behaviour of the complete protocol. Whether inferred or measured, either method of assessment is very much non-trivial. Two more of Partridge's questions are of relevance here: "Is there a formal theory for combining protocol elements?" and "Is there a theory of protocol decomposition?" [104]. In his discussion of both of them, he points to the underlying fact – that protocol composition and analysis remains very much an open-research area. In fact, of the latter question, he says "this question is, for me, the most difficult...the research space is completely empty and the results would become an integral part of how we devised protocols in the future". The challenges of addressing what represents an optimal or even sufficient solution for a scenario are as yet unsolved. Significant future work addressing these two challenges is necessary in order to make it possible to employ machine learning to move the Anatomy from a tool for evolvable design to *evolving* design.

An eighth challenge is the manner in which decoupling of functionality into elements may cause breakdown in the behaviour of some CR MAC mechanisms when split across element boundaries, where the algorithm was originally devised for a monolithic system. In the monolithic system, the algorithm may have employed a high degree of feedback between what might otherwise be parts easily separable into elements. In the Anatomy, such a mechanism may require a very large Anatomy element, or a group of elements with complex inter-element dependencies. The alternative is to remove some of the functionality, possibly reducing performance slightly. While this is a challenge to the Anatomy, it represents the challenge of modular implementations in general. The twin compromises offered by the Anatomy are firstly the use of responsibilities and awareness as grounds for the situation of elements, and secondly the refusal to prescribe a level of granularity, so that designers may maintain the monolithic nature of some CR MAC elements if they so choose.

A ninth challenge is that while the mode of data-flow between Anatomy elements has not strictly been defined, depending on the choice employed for an implementation, behaviour might be affected. Over the course of this thesis, the manner of data-flow has not been discussed. However, for the purposes of the implementations presented, it should be noted that Iris employs a block-wise data-flow model. Blocks of data are passed fully constructed to `Components`, processed, and output to the next `Component(s)` in the chain only when processing for the entire block is complete. Some CR MAC mechanisms might potentially perform actions on the basis of received data as it arrives, without having received entire blocks. They rely on stream-wise data reception. Such an element may sometimes wish to act before the course of a reception of a packet is complete. If an element relying on stream-wise data is implemented in a block-wise system, its attempts to "interrupt" a stream of data will not succeed, as the action it attempts to interrupt will already have been completed. This is not strictly a failure of the Anatomy. Rather, it is a failure introduced by an assumption of data-flow model on the part of these particular designs. Similarly, it can be assumed that certain designs might rely on a block-based data-flow. The ideal solution would be implementation on a platform capable of allowing a choice of data-flow models, with possibly both existing in a single implementation. In such a scenario, component dependency on a particular data-flow model would represent an additional characterisation to be represented on the Anatomy Schema. It should however also be noted that discussion of the data-flow model is a reflection of the non-prescriptive nature of the Anatomy. Strictly speaking, within the Open Systems Interconnect (OSI) layering concept, all inter-layer communication is packet based, divided into the Service Data Units (SDUs) received from higher layers, and following processing, output as Protocol Data Units (PDUs). Therefore block-wise data-flow is a reasonable assumption to make.

A tenth challenge is that as presented in this thesis, the Anatomy does not present knobs, meters and contracts to higher layers. As discussed in Chapter 3, §3.4.3 however, this provision is easily made and represented on an Anatomy Schema. Furthermore, it is reasonable to assume that such provisions will be desired by future cognitive networking techniques and protocols.

The eleventh and final challenge is the possibility that the Anatomy may not be capable of representing unforeseen aspects of the CR MAC design-space. While this possibility cannot be conclusively dismissed, two factors combine to make it unlikely: Firstly, the broad review of the literature that led to the creation of the Anatomy helps to ensure that it was designed to be capable of representing the widest range of protocols foreseeable given the current state-of-the-art. Secondly, the unprescriptive nature of the Anatomy – in that it does not mandate a specific layout, granularity of protocol representation, or even that every responsibility be met and role be played – lends it flexibility, which further reduces the risk. It is perhaps more likely that future increases in protocol complexity would require further subdivision of CR MAC responsibilities and awarenesses, but this would represent a requirement for further enhancement to the Anatomy, rather than heralding its obsolescence.

5.5 Comparing the Anatomy with related research

With the Anatomy of a CR MAC introduced, discussed in detail and demonstrated in action over the preceding chapters, an obvious question to ask is what research has been performed that is similar? Recall that in the literature survey in Chapter 2, §2.4.3, a select number papers concerned with the challenges of flexible MAC in general were explored. These papers rather than proposing specific protocols or mechanisms, were concerned with the greater challenges related to realising flexible and CR MACs, and aiding their design and construction, much as this thesis is. It will therefore be of interest to re-explore them in the context of the Anatomy, to examine shared similarities – as well as the points of difference – with the Anatomy and its approach.

The "TRUMP" framework proposed by Zhang *et al.* differs from the Anatomy in that it is a specific implementation mechanism, and is bound to a fixed and inflexible PHY implementation [62]. In being so, it neglects the importance of PHY dependencies. It does not discuss awareness or spectrum resource allocation versus exploitation. And furthermore, the constituent blocks of TRUMP are significantly more granular than those of the Anatomy, operating at the level of state-machine constituents. TRUMP also lacks a structured graphical representation and design approach, as allowed by the Anatomy Schema, and therefore there is no clear ground-up method for construction of a new protocol.

Notwithstanding these limitations, TRUMP could be a candidate for the implementation of an Anatomy approach. Anatomy Elements could be constructed of more granular TRUMP parts, in order to take advantage of its leveraging of Field Programmable Gate Array (FPGA) resources to provide tight coupling between the MAC and PHY, and near dedicated hardware levels of performance. As such, it represents research complimentary to the Anatomy.

The "MultiMAC" proposed by Doerr *et al.* is at odds with the approach of the Anatomy of a CR MAC, as it considers the whole-protocol level to be the best to operate at when potentially reconfiguring behaviour [64]. This thesis would disagree, noting that parametric reconfiguration of monolithic protocols can never bring the flexibility of a more granular approach. MultiMAC's strength however lies in its rule based MAC selection element, the thinking behind which might inform construction of an engine to select Anatomy elements online, to enable construct CR MAC protocols as needed.

Huang *et al.* proposed an "adaptive MAC" (AMAC) [65]. Similar to MultiMAC, it also operates at the whole-protocol level, a granularity at odds with that favoured by the Anatomy of a CR MAC. However, adoption of the global control plane as a mechanism through which reconfiguration decisions are made on a group level is an intriguing possibility for future work, representable on an Anatomy schema as an awareness element for distribution of in-system and external awareness between devices.

While suitable for MAC-layer implementations, the "Click" modular router was not designed with CR in mind [67]. Mandke *et al.* extend the framework by combining it with a GNU Radio PHY [68], but still it does not in its current form take any account of PHY dependencies. Nonetheless, it is possible that with modifications to take account of the need for a possibly non-standard PHY interface, and a careful choice of PHY to operate with, that Click could represent a possible choice of platform on which to take the Anatomy's approach to implementation.

The MetaMAC proposal of Faragó *et al.* [69] shares some of the common thinking of Huang's AMAC and Doerr's MultiMAC. As with MultiMAC and AMAC, it too operates using a library of self-contained complete protocols. It does differ slightly in that these self-contained protocols all run simultaneously, and their resulting decisions are combined to produce the ultimate per-timeslot access decision. It is not clear that this is feasible in a real-world implementation. Furthermore, though the authors claim that the model is extendible to a non time-slotted scenario, they do not clarify how this might be done. Finally, there is little discussion of PHY interactions and dependencies, or indeed of the inputs to the individual constituent MACs, which will surely be of great importance in a flexible CR MAC.

The thinking of Lichte *et al.*is similar to the thinking behind the Anatomy and its process: namely, that viewed through the appropriate lens, complex behaviour is made up of simple elements [70]. Their proposal does not represent an ideal choice for CR MAC in that it is simplistic in terms of its flexibility, and represents a mechanism for the control of a single-channel static PHY-layer in a Carrier Sense Multiple Access (CSMA)-like manner. In that sense while it is not directly competitive with the Anatomy, it is also not directly of use in its approach. It should be noted that with further exploration it could perhaps be thought of as a formal representation of the mechanism *inside* some Anatomy elements.

Sun *et al.* and Teng *et al.* both take a very different view of system behaviour in their proposals [71, 72]. As a result, both offer an outcome-oriented characterisation of a MAC protocol, whereas the Anatomy is a process-oriented characterisation.

They are orthogonal methods of analysis of the same problem.

Finally, in their state-machine based framework Xiao *et al.* also offer a representation that from a granularity point of view is similar to the TRUMP framework [73]. Indeed, their work is noted as a source of inspiration by the authors of TRUMP. There is however little in the way of PHY interaction. As a result there are limits to the range of CR MAC protocols representable with their method. Distributed multi-channel MACs in particular do not seem to fit within their framework, and the given examples are limited to centralised MACs. Furthermore, while the authors discuss implementation of real protocols, they do so solely as a conceptual exercise. No real-world implementation is performed.

Summarising, all of these works serve to advance the thinking in flexible MAC which a CR MAC will strongly benefit from incorporating. Even so, none of them serve to unify the requirements of a MAC and the PHY crossover and division between resource allocation and exploitation in the manner the Anatomy of a CR MAC does. There is no acknowledgement of the resource allocation/ exploitation divide in any of these. Furthermore, the acknowledgement of the importance of considering PHY dependencies, and of allowing control of PHY features is also lacking. And only Huang *et al.* note the importance of awareness in a CR MAC, doing so in their discussion of the global control plane feature of their AMAC proposal.

Aside from these limitations and the steps that might be needed to remedy them, the largest remaining difference between them and the view taken by the Anatomy is that of granularity. At one extreme, AMAC, MultiMAC and MetaMAC view the problems of a flexible MAC as best solved at the coarsest degree of granularity possible – that of complete protocols, deployed wholly, as needed. There is a limited degree of discussion as to allowing parametric reconfiguration of these protocols within AMAC and MultiMAC, but functionality changes of significance require a wholesale change of protocol. At the other extreme lie TRUMP, Click/ GNU Radio and the model of Xiao *et al.*, which operate at a very fine level of granularity. They construct MAC protocols of small constituent parts, at the level of state-machine elements in the case of TRUMP and Xiao's model, and a variable but usually slightly coarser level in the Click/ GNU Radioframework employed by Mandke.

As employed in the discussions earlier in this chapter, the Anatomy takes a view that lies between the two. The size of elements allows mechanisms of reasonable complexity to be embedded within them, but recombined when desired to allow for significant changes in functionality. In that sense, and as mentioned earlier in this section, the finely granular models such as TRUMP and Click could poten-

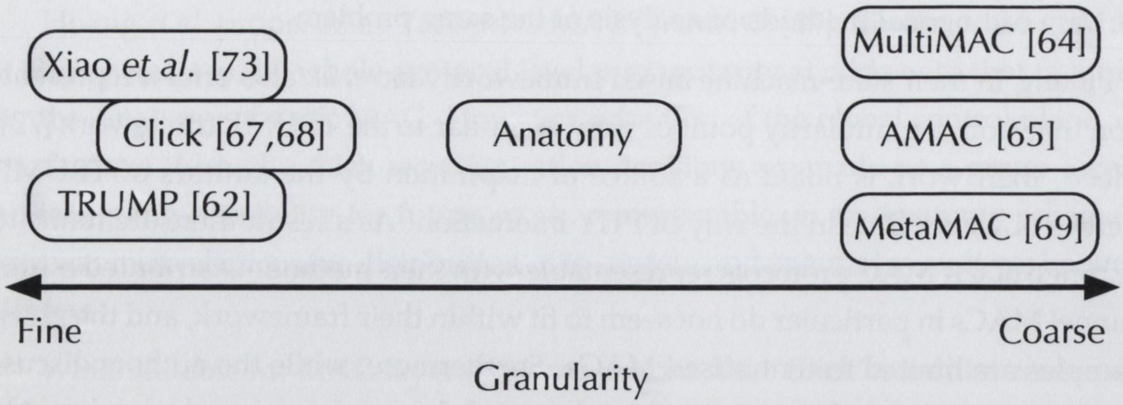

(a) The comparable flexible/ CR MAC research discussed in the section, laid in terms out of the relative granularity each solution views MACs at.

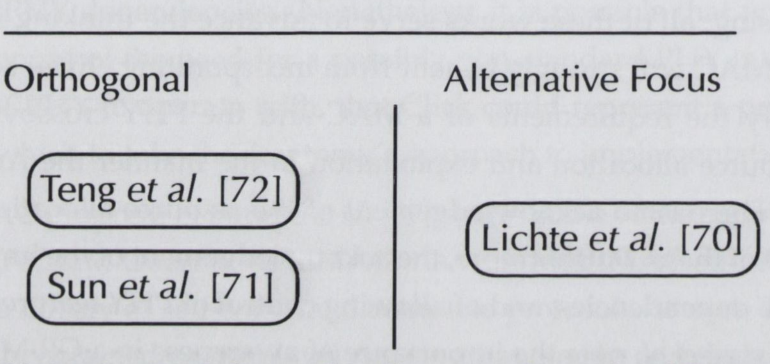

(b) Remaining flexible MAC research discussed in this section. Sun *et al.* and Teng*et al.* take an orthogonal, system-wide view, while Lichte *et al.* are concerned with describing more specific system behaviour.

Figure 5.8: Graphical comparison of flexible/ CR MAC research with the Anatomy.

tially serve as alternative candidate platforms for implementation of an Anatomy approach. Fig 5.8 lays out these proposals in accordance with their level of granularity. In the case of the Anatomy however, as was mentioned earlier in the chapter – and indeed discussed as a possible challenge to the anatomy – granularity is not strictly mandated. While the above figure depicts the granularity generally used throughout this thesis, the view used by the Anatomy may be scaled to suit the preferred granularity of a CR MAC designer. In that sense, Doerr's MultiMAC or Huang's AMAC may also be represented on a series of Anatomy Schemata, as single monolithic elements for each contained protocol. Similarly, it could be scaled to the level of TRUMP or Click. Therefore, it becomes clear that these are all compatible views of flexible or CR MACs. The view of this thesis however is that the most benefit can be reaped from the mid-level of granularity as employed throughout it. Furthermore, in addition to not prescribing a specific granularity, the Anatomy goes further than TRUMP or Click, by providing the framework on which a CR MAC's constituent parts should be laid out. It serves as the method for laying out a "recipe" for a CR MAC protocol that was noted as lacking in Chapter 2, §2.5.

5.6 Summary

To summarise, the Anatomy of a CR MAC facilitates design of CR MAC protocols with variable PHY behaviour and possibly multiple discrete PHYs. It is a method that can be applied to any specific platform of implementation. Iris is discussed in Chapter 4, §4.1 as one such platform, but simply as an example. The Anatomy requires that PHY-MAC overlap be noted in the form of PHY dependencies and contracts. It also distinguishes between the resource allocation and resource exploitation roles of a CR MAC. It varies from alternative mechanisms enabling CR and flexible MAC for these reasons, and in the flexibility of the granularity of its representation. In fact this flexibility allows the majority of research targeted at enabling CR and flexible MAC representation and implementation to be encompassed within the Anatomy, or serve as a platform for implementation of an Anatomy designed CR MAC.

further than TRUMP or Click by providing the framework on which a CR MAC's constituent parts should be laid out. It serves as the method for [illegible] for a CR MAC protocol that was noted as lacking in Chapter 2 [illegible].

To summarise, the Anatomy of a CR MAC facilitates design of CR MAC protocols with variable PHY behaviour and possibly multiple discrete PHYs. It is a method that can be applied to any specific platform of implementation. Iris is discussed in Chapter 4, §4.1 as one such platform, but simply as an example. The Anatomy requires that PHY-MAC overlap be noted in the form of PHY dependencies and contracts. It also distinguishes between the resource allocation and resource exploitation roles of a CR MAC. It varies from alternative mechanisms enabling CR and flexible MAC for these reasons, and in the flexibility of the granularity of its representation. In fact this flexibility allows the majority of research targeted at enabling CR and flexible MAC representation and implementation to be encompassed within the Anatomy, or serve as a platform for implementation of an Anatomy designed CR MAC.

(b) remaining flexible MAC research discussed in this section. Sutton et al. and Jorge et al. take an orthogonal, system-wide view, while Lahiri et al. are concerned with describing more specific system behaviour.

Figure 5.8: Graphical comparison of flexible CR MAC research with the Anatomy.

tially serve as alternative candidate platforms for implementation of an Anatomy approach. Fig 5.8 lays out these proposals in accordance with their level of granularity. In the case of the Anatomy however, as was mentioned earlier in the chapter – and indeed discussed as a possible challenge to the anatomy – granularity is not strictly mandated. While the above figure depicts the granularity generally used throughout this thesis, the view used by the Anatomy may be scaled to suit the preferred granularity of a CR MAC designer. In this sense, Doerr's MultiMAC or Huang's AMAC may also be represented on a series of Anatomy Schemata as single monolithic elements for each contained protocol. Similarly, it could be scaled to the level of TRUMP or Click. Therefore, it becomes clear that these are all compatible views of flexible or CR MACs. The view of this thesis however is that the most benefit can be reaped from the mid-level of granularity as employed throughout it. Furthermore, in addition to not prescribing a specific granularity, the Anatomy goes

6 Conclusions

This final chapter of the thesis summarises the work that has taken place and casts an eye to the future for Cognitive Radio (CR) Medium Access Control (MAC) research.

6.1 Contributions

The thesis focused on Medium Access Control protocols for Cognitive Radio. In particular it explored how to systematically define and design CR MACs protocols. The core contribution of this thesis was the introduction of the Anatomy of a CR MAC. The Anatomy provides a structured and consistent means of identifying the elements of a CR MAC in a manner that takes account of the special challenges of CR, such as the additional dependencies on a flexible physical (PHY) layer, and the extra functionality associated with the identification and sharing of spectral resources.

The Anatomy was put into action in different ways. It was used to describe and chart existing CR MAC solutions. It was used to compare different CR MACs. It was used to combine elements from various MACs to create new CR MACs, and it was used to design a specific CR MAC from scratch, dubbed "CycloMAC". CycloMAC was subsequently implemented on a platform consisting of the Iris software radio and the USRP RF-frontend.

A significant body of research has already been carried out in the area of CR MACs. As mentioned earlier, in the region of four to five hundred papers exist which explicitly self-nominate as CR MAC works, and many hundreds more cover related topics. The Anatomy of a MAC offers a mode of making sense of all of this work; a way through the literature; and as mentioned and demonstrated a number of times in this thesis, a means of making use of the material in a structured manner.

6.2 The Anatomy and the ideal CR MAC

It is worth thinking in broader terms about where CR MAC research is going. Recall that Chapter 2 framed its discussion of CR MAC, and led-in to the introduction of the Anatomy of a CR MAC by questioning the possible existence and definition of an ideal CR MAC. With the Anatomy introduced, applied, discussed and dissected, this topic may once again be returned to. That is, what is an ideal CR MAC, especially when as has been discussed it is clear that most authors say CR but mean opportunistic DSA? Chapter 2 proposed that it be defined as a protocol capable of evolving its own behaviour over time in response to stimulus. Nothing discussed in the intervening chapters has suggested a better definition. Two suggested extremes of possibility were i.) a super-protocol, capable of doing performing well or even optimally in all situations and ii.) a library of protocols, each deployed in a specific situation.

The first definition intuitively seems rather unlikely, and it is difficult to find anything in the literature attempting to claim such, even if in the conventional wireless world one particular MAC protocol – the IEEE802.11 Distributed Coordination Function (DCF) [3] – has captured a great deal of imagination. Even devising simple hypothetical wireless MAC scenarios, it is not difficult to come up with a selection which require very different performance characteristics. For example, a scenario in which a Time Division Multiple Access (TDMA) MAC performs extremely well is very different from one in which a Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) one is a good choice. Hybrid MACs do offer some solution, but at the expense of overhead and complexity, and their own rigidities and caveats.

The second definition seems one with great deal more potential. For the same reason a single protocol is unlikely to suit all situations, the ability to select a protocol from a library seems of significantly greater benefit. The question in this case is the extent to which these protocols are monolithic and/ or flexible. As seen in the previous section, this is an approach that has been adopted and explored in the literature, and some examples were provided in the last chapter, including MetaMAC and MultiMAC in §5.5.

Yet, the best option would seem to lie in the middle ground, in which protocols are constructed of smaller elements, the arrangement and possibly parameters of which can be altered over time in response to experience. The Anatomy could aid such a vision CR MAC, by serving as a framework for implementation of a MAC capable of evolving its own behaviour.

Elements could be replaced by more advanced elements – the CycloMAC **Scheduler** element might learn the best way to schedule, for example, assisted by greater internal awareness. The frequency selection process in the **Cyclostationary bootstrap, rendezvous and recovery** element might cease to be simple, perhaps choosing frequencies with knowledge of a cheap transmitter's poor filtering or non-linearities. These might result in significant out-of-band emissions, but could be taken into account in a careful decision of where to transmit, a penalty possibly offset by the reduced hardware and energy cost of the overall CR device. On this level, piece by piece, protocol behaviour could be evolved, balancing improvements and tradeoffs.

And on a grander scale, an overarching entity such as a Cognitive Engine might become responsible for the online selection and combination of Anatomy elements. Given future research to address the sixth and seventh challenges discussed in Chapter 5, §5.4, and by making use of more comprehensive learning models, it could build and test the performance of protocols. It could add to them iteratively using advanced optimisation techniques, such as genetic algorithms, and approximation techniques such as artificial neural-networks. While this may seem a far-flung possibility, the first step would be a protocol "slightly" reorganising itself, making small, single element changes of its configuration. Rather than the example above of elements learning and making decisions to improve their internal behaviour, the improvements would be made from the outside. Then as techniques grow in power – and indeed experience – further and more radical reconfigurations would become understood, and eventually, commonplace. A CR MAC would hence become a library of protocol elements, constructing itself into any temporary protocol suited to its environment.

In another sense, the recognition of the need to not over-specify and prescribe exact MAC behaviour that is implicit in the viewpoint of the Anatomy, is one that can be seen beginning to permeate real-world systems. LTE-Advanced specifications, for example, do not mandate every minute detail of the system. In certain cases, such as scheduling, there is scope to implement widely varying solutions [105]. Its carrier aggregation function too leaves tremendous scope for flexibility, both at implementation and run-time, in terms of how the resources of a network are distributed to user equipment [106]. This move to prioritising task completion over the method used to do so is an important aspect of CR, that appears often forgotten – there are a great many possible situations in which terrible performance is quite optimal. The growing acceptance of the merit of such sparse specification silently echoes the thrust of this thesis, that a single CR MAC protocol is not enough: A protocol must be constructed as suited to its device's environment.

Current interpretations of CR as opportunistic Dynamic Spectrum Access (DSA) mean that as yet, future implementations remain fixed. They will be complex MAC protocols, certainly, but remain fully specced at design time, and will not change afterwards in unforeseen ways. Whatever the path the future of CR MAC takes, such "truer" CRs will undoubtedly require more flexible MACs, at which point a tool such as the Anatomy will go from highly useful, to necessary.

6.3 Future work

This thesis introduced the Anatomy of a CR MAC, and demonstrated its use as a tool for implementation with several example CR MACs, as well as "off-line" to examine numerous protocols in the literature. Their remain however, several avenues for future research, outlined here.

6.3.1 Higher-layer interactions

There is scope for research into the consideration of contracts with higher-layers, by means of implementation and experimentation. Existing higher layer implementations at CTVR Telecommunications Research Centre (CTVR) could be leveraged to consider the benefits of CR MAC provided MAC contracts and dependencies to Mobile Ad-Hoc NETwork (MANET) routing-protocols.

6.3.2 Implementation on other Software Defined Radio (SDR) platforms

The Anatomy was used to guide implementation here using the Iris SDR platform. Deployment of the same methods on other example platforms would allow exploration of any enhancements necessary to ensure a completely generalised model for CR MACs. It would also be of particular interest to investigate how use of a stream-based data-passing architecture might change inter-element interactions, and perhaps change or negate the assistance that can be provided by abstraction elements.

6.3.3 Implementation with greater PHY-dependencies

Another avenue for future research would be implementations taking advantage of more advanced hardware. Features of the latest generations of USRP hardware such as access to significant idle FPGA fabric could be employed to implement interesting PHY-dependent features for a CR MAC to control. And alternative hardware with closer frontend-GPP integration would allow for significant performance improvements in speed, to address some of the problems relating to the value of experimentation discussed in §4.2.5 of chapter 4.

6.3.4 Improved usability

From a usability point-of-view, the current implementation of the Anatomy in Iris remains a command line application. A Graphical User Interface (GUI) allowing construction of protocols on a Schema would be of use as a learning tool, particularly to new users of the Anatomy, and researches new to CR MAC in general. Implementing such a tool would also require the ability to rapidly verify input, output and dependency satisfaction of all elements.

6.3.5 CycloMAC enhancements

Finally, in the context of the implemented CycloMAC protocol, there is scope for a large number of improvements to be made. It should be noted however, that minor static improvements for their own sake would represent a direct contradiction of the spirit of this thesis.

6.4 Final word

A CR MAC may one day be one that takes account of its environment to change its behaviour, and may react to the same situation in many different ways, on account of learning and perhaps its own experimentation. But truthfully speaking, what "it" is, will be ever-changing; will be ever reacting; will at no point be fixed; and as such, there will be no ideal CR MAC.

Appendices

A CycloMAC Iris xml configuration

A.1 CycloMAC

The xml below represents an Iris configuration for a master device using the implemented "CycloMAC" protocol, as discussed in Chapter 4, §4.2.

```xml
<?xml version="1.0" encoding="utf-8" ?>

<softwareradio name="AnatomyCyclo0">

  <controller class="synchronisation" />
  <controller class="usrptime" />
  <controller class="anatomysignature" />

  <engine name="stackengine1" class="stackengine">

    <component name="filewriter1" class="filewriter">
      <parameter name="filename" value="output.txt"/>
      <port name="topport1" class="io"/>
      <port name="bottomport1" class="io"/>
    </component>

    <component name="filereader1" class="filereader">
      <parameter name="filename" value="node0_data.bin"/>
      <parameter name="delay" value="20"/>
      <parameter name="blocksize" value="734"/>
      <port name="topport1" class="io"/>
      <port name="bottomport1" class="io"/>
    </component>

    <component name="arq1" class="arq">
      <parameter name="timeout" value="50"/>
      <parameter name="ownaddress" value="000000000000"/>
      <parameter name="destinationaddress" value="ffffffffffff"/>
      <parameter name="maxbackoff" value="2"/>
      <parameter name="debug" value="false"/>
      <port name="topinputport" class="io"/>
      <port name="topoutputport" class="io"/>
      <port name="dataport" class="io"/>
      <port name="controlport" class="io"/>
      </component>

```

```
    <component name="masterexpcontrol" class="expcontrolchannel">
      <port name="topport1" class="io"/>
      <port name="bottomport1" class="io"/>
    </component>

    <component name="datachannel1" class="datachannel">
      <port name="topport1" class="io"/>
      <port name="bottomport1" class="io"/>
    </component>

    <component name="scheduler1" class="scheduler">
      <parameter name="debug" value="false"/>
      <parameter name="beaconlength" value="1"/>
      <parameter name="expcontrolslotlength" value="1"/>
      <parameter name="dataslotlength" value="10"/>
      <parameter name="maxexpcontrolbytes" value="40"/>
      <parameter name="maxdatabytes" value="750"/>
      <port name="expcontrolport" class="io"/>
      <port name="dataport" class="io"/>
      <port name="bottomport" class="io"/>
    </component>

    <component name="sync1" class="clustersync">
      <parameter name="ownaddress" value="000000000000"/>
      <parameter name="debug" value="true"/>
      <parameter name="cooldowntime" value="500"/>
      <parameter name="maintenanceinterval" value="1000"/>
      <parameter name="synctimeout" value="100"/>
      <parameter name="linktimeout" value="5000"/>
      <port name="topport" class="io"/>
      <port name="bottominputport" class="io"/>
      <port name="bottomoutputport" class="io"/>
    </component>

  </engine>

  <engine name="pnenginerx" class="pnengine">

    <component name="ofdmdemod1" class="simpleofdmdemod">
      <port name="input1" class="input"/>
      <port name="output1" class="output"/>
      </component>

    <component name="sigdetect1" class="cyclosignaturedetector">
      <parameter name="debug" value="false"/>
      <parameter name="fftsize" value="512"/>
      <parameter name="cyclicfrequency" value="62500"/>
      <parameter name="bandwidth" value="1000000"/>
      <parameter name="numwindows" value="32"/>
      <parameter name="freqsmoothingorder" value="16"/>
      <parameter name="freqshift" value="-10"/>
      <parameter name="detectthreshold" value="0.4"/>
      <parameter name="triggerreconfig" value="false"/>
      <parameter name="passthroughdata" value="false"/>
      <parameter name="freqacquisition" value="true"/>
      <parameter name="settlingtime" value="2"/>
      <parameter name="minfrequency" value="5000000000"/>
      <parameter name="maxfrequency" value="5005000000"/>
```

```
      <port name="input1" class="input"/>
      <port name="output1" class="output"/>
    </component>

    <component name="uhdrx1" class="usrpuhdrx">
      <parameter name="args" value=""/>
      <parameter name="rate" value="1000000"/>
      <parameter name="frequency" value="5003000000"/>
      <parameter name="gain" value="40"/>
      <parameter name="outputblocksize" value="4096"/>
      <parameter name="fixlooffset" value="5000000"/>
      <port name="output1" class="output"/>
      </component>

  </engine>

  <engine name="pnenginetx" class="pnengine">


    <component name="ofdmmod1" class="simpleofdmmod">
      <parameter name="embedsignature" value="true"/>
      <parameter name="numfftbins" value="512"/>
      <parameter name="bitsperdatasymbol" value="1"/>
      <parameter name="enable" value="false"/>
      <port name="input1" class="input"/>
      <port name="output1" class="output"/>
    </component>

    <component name="scaler1" class="signalscaler">
      <parameter name="maximum" value="0.8"/>
      <port name="input1" class="input"/>
      <port name="output1" class="output"/>
    </component>

    <component name="uhdtx1" class="usrpuhdtx">
      <parameter name="args" value=""/>
      <parameter name="rate" value="2000000"/>
      <parameter name="frequency" value="5003000000"/>
      <parameter name="gain" value="0"/>
      <parameter name="streaming" value="false"/>
      <parameter name="fixlooffset" value="5000000"/>
      <parameter name="checkerrors" value="true"/>
      <port name="input1" class="input"/>
    </component>

  </engine>
  <link source="filewriter1.bottomport1" sink=                    ▼
     "arq1.topoutputport" />
  <link source="filereader1.bottomport1" sink=                    ▼
     "arq1.topinputport" />
  <link source="arq1.dataport" sink=                              ▼
     "datachannel1.topport1" />
  <link source="arq1.controlport" sink=                           ▼
     "expcontrol1.topport1" />
  <link source="datachannel1.bottomport1" sink=                   ▼
     "scheduler1.dataport" />
  <link source="expcontrol1.bottomport1" sink=                    ▼
     "scheduler1.expcontrolport" />
  <link source="scheduler1.bottomport" sink=                      ▼
     "sync1.topport" />
```

```xml
  <!--Tx chain-->
  <link source="sync1.bottomoutputport" sink="ofdmmod1.input1" />
  <link source="ofdmmod1.output1" sink="scaler1.input1" />
  <link source="scaler1.output1" sink="uhdtx1.input1" />
  <!--Rx chain-->
  <link source="uhdrx1.output1" sink="sigdetect1.input1" />
  <link source="sigdetect1.output1" sink="ofdmdemod1.input1" />
  <link source="ofdmdemod1.output1" sink="sync1.bottominputport" />
</softwareradio>
```

Acronyms

ACK	Acknowledgment packet
ACM	Association for Computing Machinery
AI	Artificial Intelligence
AP	Access Point
ARQ	Automatic Repeat reQuest
ASIC	Application Specific Integrated Circuit
BER	Bit Error Rate
CCC	Common Control Channel
CDMA	Code Division Multiple Access
COTS	Consumer Off-The-Shelf
CR	Cognitive Radio
CRC	Cyclic Redundancy Check
CSMA	Carrier Sense Multiple Access
CSMA/CA	Carrier Sense Multiple Access with Collision Avoidance
CTS	Clear To Send packet
CTVR	CTVR Telecommunications Research Centre
DCF	Distributed Coordination Function
DFT	Discrete Fourier Transform
DSA	Dynamic Spectrum Access
DSP	Digital Signal Processing
DSSS	Direct Sequence Spread Spectrum
EC	Edge-case
FCC	Federal Communications Commission

FDD	Frequency Division Duplexing
FDMA	Frequency Division Multiple Access
FEC	Forward Error Correction
FFT	Fast Fourier Transfrom, a fast implementation of the DFT
FHSS	Frequency Hopping Spread Spectrum
FPGA	Field Programmable Gate Array
GPP	General Purpose Processor
GPRS	General Packet Radio Service
GRAP	Group Randomly Addressed Polling
GSM	Global System for Mobile communications
GUI	Graphical User Interface
IEEE	Institute of Electrical and Electronics Engineers
ISM	Industrial Scientific & Medical band, a group of frequencies set aside in most countries in which unlicensed use is permitted subject to regulator mandated power constraints
LAN	Local Area Network
LTE	Long-Term Evolution
MAC	Medium Access Control, a sub-layer of the OSI stack model
MAI	Multiple Access Interference
MANET	Mobile Ad-Hoc NETwork
NAV	Network Allocation Vector, a counter indicating expected transaction time on-the-air included in control packets in some MAC protocols.
NC-OFDM	Non-Contiguous Orthogonal Frequency Division Multiplexing
OFDM	Orthogonal Frequency Division Multiplexing
OSI	Open Systems Interconnect
PRMA	Packet Reservation Multiple Access
PDU	Protocol Data Unit
PHY	physical layer, a layer of the OSI stack model
PU	Primary User
RF	radio frequency
RTS	Ready To Send packet
SCD	Single Collision Domain

SDMA	Space Division Multiple Access
SDR	Software Defined Radio
SDU	Service Data Unit
SINR	Signal to Interference plus Noise Ratio
SNR	Signal to Noise Ratio
SP	Split-Phase
SU	Secondary User
TDD	Time Division Duplexing
TDMA	Time Division Multiple Access
THSS	Time-Hopping Spread Spectrum
UMTS	Universal Mobile Telecommunications System
USRP	Universal Software Radio Peripheral
WARP	Wireless Open-Access Research Platform, a Software Defined Radio (SDR) platform for experimentation created at Rice University

FDD	Frequency Division Duplexing
FDMA	Frequency Division Multiple Access
FEC	Forward Error Correction
FFT	Fast Fourier Transform
FHSS	Frequency Hopping Spread Spectrum
FPGA	Field Programmable Gate Array
GPP	General Purpose Processor
GPRS	General Packet Radio Service
GRAP	Group Randomly Addressed Polling
GSM	Global System for Mobile Communications
GUI	Graphical User Interface
IEEE	Institute of Electrical and Electronics Engineers
ISM	[illegible] in most countries in which unlicensed use is permitted subject to regulator mandated power constraints
LAN	Local Area Network
LTE	Long Term Evolution
MAC	Medium Access Control, a sub-layer of the OSI stack model
MAI	Multiple Access Interference
MANET	Mobile Ad-Hoc NETwork
NAV	Network Allocation Vector, a counter indicating expected transaction time on-the-air included in control packets in some MAC protocols
NC-OFDM	Non Contiguous Orthogonal Frequency Division Multiplexing
OFDM	Orthogonal Frequency Division Multiplexing
OSI	Open Systems Interconnect
PRMA	Packet Reservation Multiple Access
PDU	Protocol Data Unit
PHY	physical layer, a layer of the OSI stack model
PU	Primary User
RF	radio frequency
RTS	Ready To Send packet
SCD	Single Collision Domain

Bibliography

[1] J. Mitola and G. Maguire, "Cognitive Radio: Making Software Radios More Personal," *IEEE Personal Communications*, vol. 6, no. 4, pp. 13–18, August 1999. [Online]. Available: 10.1109/98.788210

[2] J. Mitola III, "Cognitive Radio: An integrated agent architecture for software defined radio," Ph.D. dissertation, Royal Institute of Technology, Stockholm, Sweden, May 2000.

[3] "IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,," *IEEE Std 802.11-2007 (Revision of IEEE Std 802.11-1999)*, pp. 1–1076, June 2007. [Online]. Available: 10.1109/IEEESTD.2007.373646

[4] "IEEE Standard for Local and metropolitan area networks Part 16: Air Interface for Broadband Wireless Access Systems," *IEEE Std 802.16-2009 (Revision of IEEE Std 802.16-2004)*, pp. C1–2004, May 2009. [Online]. Available: 10.1109/IEEESTD.2009.5062485

[5] R. M. Brown, *Sudden Death*. Bantam Books, 1983, p. 68.

[6] T. W. Rondeau, "Application of Artificial Intelligence to Wireless Communications," Ph.D. dissertation, Virginia Polytechnic Institute and State University, September 2007.

[7] L. E. Doyle, *Essentials of Cognitive Radio*. Cambridge University Press, 2009.

[8] T. Yucek and H. Arslan, "A survey of spectrum sensing algorithms for cognitive radio applications," *IEEE Communications Surveys & Tutorials*, vol. 11, no. 1, pp. 116–130, January 2009. [Online]. Available: 10.1109/SURV.2009.090109

[9] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty, "NeXt generation/-dynamic spectrum access/cognitive radio wireless networks: A survey," *Computer Networks*, vol. 50, no. 13, pp. 2127–2159, September 2006. [Online]. Available: 10.1016/j.comnet.2006.05.001

[10] W. Lehr and J. Crowcroft, "Managing shared access to a spectrum commons," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, November 2005, pp. 420–444. [Online]. Available: 10.1109/DYSPAN.2005.1542658

[11] J. Haartson and S. Mattisson, "Bluetooth Core Specification," Bluetooth SIG, Standard, 1998.

[12] W. W. Peterson and D. T. Brown, "Cyclic Codes for Error Detection," *Proceedings of the IRE*, vol. 49, no. 1, pp. 228–235, January 1961. [Online]. Available: 10.1109/JRPROC.1961.287814

[13] N. Abramson, "THE ALOHA SYSTEM: another alternative for computer communications," in *Proceedings of the '70 AFIPS Fall Joint Computer Conference*, November 1970, pp. 281–285. [Online]. Available: 10.1145/1478462.1478502

[14] L. Kleinrock and F. Tobagi, "Packet Switching in Radio Channels: Part I–Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics," *IEEE Transactions on Communications*, vol. 23, no. 12, pp. 1400–1416, December 1975. [Online]. Available: 10.1109/TCOM.1975.1092768

[15] A. C. V. Gummalla and J. O. Limb, "Wireless medium access control protocols," *IEEE Communications Surveys & Tutorials*, vol. 3, no. 2, pp. 2–15, April 2000. [Online]. Available: 10.1109/COMST.2000.5340799

[16] F. 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*, vol. 23, no. 12, pp. 1417–1433, December 1975. [Online]. Available: 10.1109/TCOM.1975.1092767

[17] P. Karn, "MACA – A new channel access method for packet radio," in *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, September 1990, pp. 134–140.

[18] M. Z. Brodsky and R. T. Morris, "In defense of wireless carrier sense," *SIGCOMM Computer Communications Review*, vol. 39, no. 4, pp. 147–158, October 2009. [Online]. Available: 10.1145/1594977.1592587

[19] R. Dixon, *Spread spectrum systems*. John Wiley & Sons, Inc., 1976.

[20] A. J. Viterbi, *CDMA: Principles of Spread Spectrum Communication*. Addison-Wesley, 1995.

[21] A. K. L. Yau, P. Komisarczuk, and P. D. Teal, "On Multi-Channel MAC Protocols in Cognitive Radio Networks," in *Australasian Telecommunication Networks and Applications Conference (ATNAC)*, December 2008, pp. 300–305. [Online]. Available: 10.1109/ATNAC.2008.4783340

[22] Z. J. Haas and J. Deng, "Dual busy tone multiple access (DBTMA) – a multiple access control scheme for ad hoc networks," *IEEE Transactions on Communications*, vol. 50, no. 6, pp. 975–985, June 2002. [Online]. Available: 10.1109/TCOMM.2002.1010617

[23] V. Shrivastava, S. Rayanchu, J. Yoonj, and S. Banerjee, "802.11n under the microscope," in *ACM SIGCOMM Conference on Internet Measurement*, October 2008, pp. 105–110. [Online]. Available: 10.1145/1452520.1452533

[24] D. J. Goodman, R. A. Valenzuela, K. T. Gayliard, and B. Ramamurthi, "Packet Reservation Multiple Access for Local Wireless Communications," *IEEE Transactions on Communications*, vol. 37, no. 8, pp. 885–890, August 1989. [Online]. Available: 10.1109/26.31190

[25] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving Single Channel, Full Duplex Wireless Communication," in *ACM International Conference on Mobile Computing and Networking (MOBICOM)*, September 2010, pp. 1–12. [Online]. Available: 10.1145/1859995.1859997

[26] M. Duarte and A. Sabharwal, "Full-Duplex Wireless Communications Using Off-The-Shelf Radios: Feasibility and First Results," in *Asilomar Conference on Signals, Systems and Computers*, November 2010, pp. 1558–1562. [Online]. Available: 10.1109/ACSSC.2010.5757799

[27] S. Sen, R. Roy Choudhury, and S. Nelakuditi, "CSMA/CN: Carrier Sense Multiple Access with Collision Notification," in *ACM International Conference on Mobile Computing and Networking (MOBICOM)*, September 2010, pp. 25–36. [Online]. Available: 10.1145/1859995.1859999

[28] ETSI, "ETSI TR 102 683 V1.1.1 (2009-09) Reconfigurable Radio Systems (RRS); Cognitive Pilot Channel (CPC)," European Telecommunications Standards Institute, Tech. Rep., 2009.

[29] M. Felegyhazi, M. Cagalj, and J. P. Hubaux, "Efficient MAC in cognitive radio systems: A game-theoretic approach," *IEEE Transactions on Wireless Communications*, vol. 8, no. 4, pp. 1984–1995, April 2009. [Online]. Available: 10.1109/TWC.2009.080284

[30] V. B. Mišić and J. Mišić, "Performance of a Cognitive MAC Protocol with Transmission Tax," *IEEE Conference on Computer Communications (INFOCOM) – Workshops*, pp. 1–5, March 2010. [Online]. Available: 10.1109/INFCOMW.2010.5466715

[31] C. Cordeiro and K. Challapali, "C-MAC: A Cognitive MAC Protocol for Multi-Channel Wireless Networks," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, April 2007, pp. 147–157. [Online]. Available: 10.1109/DYSPAN.2007.27

[32] S. Kurkowski, T. Camp, and M. Colagrosso, "MANET Simulation Studies: The Incredibles," *SIGMOBILE Mobile Computer Communications Review*, vol. 9, no. 4, pp. 50–61, October 2005. [Online]. Available: 10.1145/1096166.1096174

[33] L. Lai, H. Jiang, and H. V. Poor, "Medium Access in Cognitive Radio Networks: A Competitive Multi-Armed Bandit Framework," *Asilomar Conference on Signals, Systems and Computers*, pp. 98–102, October 2008. [Online]. Available: 10.1109/ACSSC.2008.5074370

[34] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized cognitive MAC for opportunistic spectrum access in ad hoc networks: A POMDP framework," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 3, pp. 589–600, April 2007.

[35] J. Tang, S. Misra, and G. Xue, "Joint spectrum allocation and scheduling for fair spectrum sharing in cognitive radio wireless networks," *Computer Networks*, vol. 52, no. 11, pp. 2148–2158, August 2008. [Online]. Available: 10.1016/j.comnet.2008.03.010

[36] S.-J. Yoo, H. Nan, and T.-I. Hyon, "DCR-MAC: Distributed cognitive radio MAC protocol for wireless ad hoc networks," *Wireless Communications Mobile Computing*, vol. 9, no. 5, pp. 631–653, May 2009. [Online]. Available: 10.1002/wcm.v9:5

[37] Y. Ke and S.-J. Yoo, "MCR-MAC: Multi-channel cognitive radio MAC protocol for cooperative incumbent system protection in wireless ad-hoc network," *International Conference on Ubiquitous and Future Networks (ICUFN)*, pp. 6–11, June 2009. [Online]. Available: 10.1109/ICUFN.2009.5174275

[38] J. Jia, Q. Zhang, and X. Shen, "HC-MAC: A Hardware-Constrained Cognitive MAC for Efficient Spectrum Management," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 1, pp. 106–117, January 2008. [Online]. Available: 10.1109/JSAC.2008.080110

[39] L. Ma, X. Han, and C. C. Shen, "Dynamic open spectrum sharing MAC protocol for wireless ad hoc networks," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, November 2005, pp. 203–213. [Online]. Available: 10.1109/DYSPAN.2005.1542636

[40] C. Zou and C. Chigan, "On game theoretic DSA-driven MAC for cognitive radio networks," *Computer Communications*, vol. 32, no. 18, pp. 1944–1954, December 2009. [Online]. Available: 10.1016/j.comcom.2009.07.007

[41] N. Choi, M. Patel, and S. Venkatesan, "A Full Duplex Multi-channel MAC Protocol for Multi-hop Cognitive Radio Networks," in *International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM)*, June 2006, pp. 1–5. [Online]. Available: 10.1109/CROWNCOM.2006.363465

[42] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, "White space networking with Wi-Fi like connectivity," in *ACM SIGCOMM Conference on Data Communication*, August 2009, pp. 27–38. [Online]. Available: 10.1145/1592568.1592573

[43] H. A. B. Salameh, M. M. Krunz, and O. Younis, "MAC Protocol for Opportunistic Cognitive Radio Networks with Soft Guarantees," *IEEE Transactions on Mobile Computing*, vol. 8, no. 10, pp. 1339–1352, October 2009. [Online]. Available: 10.1109/TMC.2009.19

[44] H. Su and X. Zhang, "Cross-Layer Based Opportunistic MAC Protocols for QoS Provisionings Over Cognitive Radio Wireless Networks," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 1, pp. 118–129, January 2008. [Online]. Available: 10.1109/JSAC.2008.080111

[45] L. Zhai, K. Liu, Y. Liu, M. Yang, and L. Zhuang, "A Slot-Based MAC Protocol in Cognitive Radio Wireless Networks," in *International Conference on Wireless Communications, Networking and Mobile Computing (WiCOM)*, October 2008, pp. 1–4. [Online]. Available: 10.1109/WiCom.2008.292

[46] L.-C. Wang, A. Chen, and D. S. L. Wei, "A cognitive MAC protocol for QoS provisioning in ad hoc networks," *Physical Communication*, vol. 3, no. 2, pp. 105–118, June 2010. [Online]. Available: 10.1016/j.phycom.2010.01.001

[47] H. Song and X. Lin, "A Leasing Oriented MAC Protocol for High Spectrum Usage in Cognitive Radio Networks," in *IEEE Conference on Wireless and Mobile Computing, Networking and Communications (WIMOB)*, October 2009, pp. 173–178. [Online]. Available: 10.1109/WiMob.2009.38

[48] Y. Song-song and W. Wei, "A MAC Protocol for Congested Cognitive Radio Network," *Conference on Intelligent Networks and Intelligent Systems (ICINIS)*, pp. 228–232, November 2009. [Online]. Available: 10.1109/ICINIS.2009.65

[49] J. Mišić and V. B. Mišić, "Trading Delay for Fairness in MAC Design for Cognitive Personal Area Networks," *IEEE Global Telecommunications Conference (GLOBECOM)*, pp. 1–6, November 2009. [Online]. Available: 10.1109/GLOCOM.2009.5425233

[50] S. R. Singh, B. De Silva, T. Luo, and M. Motani, "Dynamic Spectrum Cognitive MAC (DySCO-MAC) for Wireless Mesh & Ad Hoc Networks," *IEEE Conference on Computer Communications (INFOCOM) – Workshops*, pp. 1–6, March 2010. [Online]. Available: 10.1109/INFCOMW.2010.5466712

[51] L. Cao and H. Zheng, "Distributed spectrum allocation via local bargaining," in *IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks (SECON)*, September 2005, pp. 475–486. [Online]. Available: 10.1109/SAHCN.2005.1557100

[52] C. N. Ververidis, J. Nasreddine, and P. Mähönen, "Distributed channel selection protocol for single-hop transmissions in cognitive radio networks," in *ACM symposium on Performance Evaluation of Wireless Ad hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, October 2009, pp. 155–156. [Online]. Available: 10.1145/1641876.1641907

[53] L. Le and E. Hossain, "OSA-MAC: A MAC Protocol for Opportunistic Spectrum Access in Cognitive Radio Networks," *IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 1426–1430, March 2008. [Online]. Available: 10.1109/WCNC.2008.256

[54] S. C. Jha, M. M. Rashid, V. K. Bhargava, and C. Despins, "OMC-MAC: An Opportunistic Multichannel MAC for Cognitive Radio Networks," *IEEE Vehicular Technology Conference (VTC 2009-Fall)*, pp. 1–5, September 2009. [Online]. Available: 10.1109/VETECF.2009.5379053

[55] L. Jiao and F. Y. Li, "MAC strategies for single rendezvous Multi-hop Cognitive Radio Networks," *IEEE Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pp. 608–612, September 2009. [Online]. Available: 10.1109/PIMRC.2009.5449836

[56] K. Ghaboosi, M. Latva-aho, and Y. Xiao, "A Distributed Multi-channel Cognitive MAC Protocol for IEEE 802.11s Wireless Mesh Networks," in *International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM)*, May 2008, pp. 1–8. [Online]. Available: 10.1109/CROWNCOM.2008.4562477

[57] D. Hu and S. Mao, "Design and Analysis of a Sensing Error-Aware MAC Protocol for Cognitive Radio Networks," in *IEEE Global Telecommunications Conference (GLOBECOM)*, November 2009, pp. 1–6. [Online]. Available: 10.1109/GLOCOM.2009.5425580

[58] J. Xiang, Y. Zhang, and T. Skeie, "Medium access control protocols in cognitive radio networks," *Wireless Communications and Mobile Computing*, vol. 10, no. 1, pp. 31–49, December 2010. [Online]. Available: 10.1002/wcm.906

[59] T. V. Krishna and A. Das, "A survey on MAC protocols in OSA networks," *Computer Networks*, vol. 53, no. 9, pp. 1377 – 1394, January 2009. [Online]. Available: 10.1016/j.comnet.2009.01.003

[60] C. Cormio and K. R. Chowdhury, "A survey on MAC protocols for cognitive radio networks," *Ad Hoc Networks*, vol. 7, no. 7, pp. 1315 – 1329, January 2009. [Online]. Available: 10.1016/j.adhoc.2009.01.002

[61] I. F. Akyildiz, W.-Y. Lee, and K. R. Chowdhury, "CRAHNs: Cognitive radio ad hoc networks," *Ad Hoc Networks*, vol. 7, no. 5, pp. 810–836, July 2009. [Online]. Available: 10.1016/j.adhoc.2009.01.001

[62] X. Zhang, J. Ansari, G. Yang, and P. Mähönen, "TRUMP: Supporting Efficient Realization of Protocols for Cognitive Radio Networks," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*. IEEE, May 2011, pp. 476–487. [Online]. Available: 10.1109/DYSPAN.2011.5936237

[63] A. Khattab, J. Camp, C. Hunter, P. Murphy, A. Sabharwal, and E. W. Knightly, "WARP: A Flexible Platform for Clean-Slate Wireless Medium Access Protocol Design," *SIGMOBILE Mobile Computer Communications Review*, vol. 12, no. 1, pp. 56–58, January 2008. [Online]. Available: 10.1145/1374512.1374532

[64] C. Doerr, M. Neufeld, J. Fifield, T. Weingart, D. C. Sicker, and D. Grunwald, "MultiMAC - an adaptive MAC framework for dynamic radio networking," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, November 2005, pp. 548–555. [Online]. Available: 10.1109/DYSPAN.2005.1542668

[65] K.-C. Huang, X. Jing, and D. Raychaudhuri, "MAC Protocol Adaptation in Cognitive Radio Networks: An Experimental Study," *Internatonal Conference on Computer Communications and Networks (ICCCN)*, pp. 1–6, August 2009. [Online]. Available: 10.1109/ICCCN.2009.5235370

[66] X. Jing and D. Raychaudhuri, "Global Control Plane Architecture for Cognitive Radio Networks," *IEEE International Conference on Communications (ICC)*, pp. 6466–6470, June 2007. [Online]. Available: 10.1109/ICC.2007.1069

[67] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The Click Modular Router," *ACM Transactions on Computer Systems*, vol. 18, pp. 263–297, August 2000. [Online]. Available: 10.1145/354871.354874

[68] K. Mandke, S.-H. Choi, G. Kim, R. Grant, R. C. Daniels, W. Kim, R. W. Heath, and S. M. Nettles, "Early Results on Hydra: A Flexible MAC/PHY Multihop Testbed," in *IEEE Vehicular Technology Conference (VTC 2007-Spring)*, April 2007, pp. 1896–1900. [Online]. Available: 10.1109/VETECS.2007.393

[69] A. Faragó, A. D. Myers, V. R. Syrotiuk, and G. V. Zaruba, "Meta-MAC protocols: automatic combination of MAC protocols to optimize performance for unknown conditions," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 9, pp. 1670–1681, September 2000. [Online]. Available: 10.1109/49.872955

[70] H. S. Lichte, S. Valentin, and H. Karl, "Automated Development of Cooperative MAC Protocols," *Mobile Network Applications*, vol. 15, pp. 769–785, December 2010. [Online]. Available: 10.1007/s11036-009-0210-5

[71] Y.-K. Sun, K.-C. Chen, and D.-C. Twu, "Generalized tree multiple access protocols in packet switching networks," in *IEEE Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, vol. 3, September 1997, pp. 918–922. [Online]. Available: 10.1109/PIMRC.1997.627020

[72] C.-M. Teng and K.-C. Chen, "A unified algorithm for wireless MAC protocols," *IEEE Vehicular Technology Conference (VTC 2002-Spring)*, vol. 1, pp. 394–398, May 2002. [Online]. Available: 10.1109/VTC.2002.1002736

[73] Z. Xiao, T. S. Randhawa, and R. H. S. Hardy, "A state-machine based design of adaptive wireless MAC layer," *IEEE Vehicular Technology Conference (VTC 2003-Spring)*, vol. 4, pp. 2837–2841, April 2003. [Online]. Available: 10.1109/VETECS.2003.1208911

[74] The MadWifi Project. [Online]. Available: http://madwifi-project.org

[75] R. Dhar, G. George, A. Malani, and P. Steenkiste, "Supporting Integrated MAC and PHY Software Development for the USRP SDR," *IEEE Workshop on Networking Technologies for Software Defined Radio Networks*, pp. 68–77, September 2006. [Online]. Available: 10.1109/SDR.2006.4286328

[76] T. Schmid, O. Sekkat, and M. B. Srivastava, "An experimental study of network performance impact of increased latency in software defined radios," in *ACM workshop on Wireless Network Testbeds, Experimental evaluation and Characterization (WinTECH)*, September 2007, pp. 59–66. [Online]. Available: 10.1145/1287767.1287779

[77] G. Nychis, T. Hottelier, Z. Yang, S. Seshan, and P. Steenkiste, "Enabling MAC protocol implementations on software-defined radios," in *USENIX symposium on Networked Systems Design and Implementation (NSDI)*, April 2009, pp. 91–105.

[78] L. A. DaSilva, A. B. MacKenzie, C. da Silva, and R. W. Thomas, "Requirements of an Open Platform for Cognitive Networks Experiments," *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pp. 1–8, October 2008. [Online]. Available: 10.1109/DYSPAN.2008.49

[79] J. Lotze, "Enabling Cognitive Radio Experimental Research," Ph.D. dissertation, University of Dublin, Trinity College, 2011.

[80] A. De Domenico, E. Strinati, and M. Benedetto, "A Survey on MAC Strategies for Cognitive Radio Networks," *Communications Surveys & Tutorials, IEEE*, vol. 14, no. 1, pp. 21–44, January 2012. [Online]. Available: 10.1109/SURV.2011.111510.00108

[81] K. Mandke, R. C. Daniels, S.-H. Choi, S. M. Nettles, and R. W. Heath, "Physical concerns for cross-layer prototyping and wireless network experimentation," in *ACM International workshop on Wireless Network Testbeds, Experimental evaluation and Characterization (WinTECH)*, September 2007, pp. 11–18. [Online]. Available: 10.1145/1287767.1287771

[82] J. Y. Yu and P. H. J. Chong, "A Survey of Clustering Schemes for Mobile Ad Hoc Networks," *IEEE Communications Surveys & Tutorials*, vol. 7, no. 1, pp. 32–48, January 2005. [Online]. Available: 10.1109/COMST.2005.1423333

[83] P. D. Sutton, J. Lötze, H. Lahlou, S. A. Fahmy, K. E. Nolan, B. Özgül, T. W. Rondeau, J. Noguera, and L. E. Doyle, "IRIS: An Architecture for Cognitive Radio Networking Testbeds," *IEEE Communications Magazine*, vol. 48, no. 9, pp. 114–122, September 2010. [Online]. Available: 10.1109/MCOM.2010.5560595

[84] Ettus Research LLC. [Online]. Available: http://www.ettus.com/

[85] Free Software Foundation, Inc. GNU Radio — The GNU Software Radio. [Online]. Available: http://www.gnu.org/software/gnuradio/

[86] C. Hunter, J. Camp, P. Murphy, A. Sabharwal, and C. Dick, "A Flexible Framework for Wireless Medium Access Protocols," *Asilomar Conference on Signals, Systems and Computers*, pp. 2046–2050, October 2006. [Online]. Available: 10.1109/ACSSC.2006.355126

[87] S. Gupta, C. Hunter, P. Murphy, and A. Sabharwal, "WARPnet: clean slate research on deployed wireless networks," in *ACM symposium on Mobile ad hoc networking and computing (MobiHoc)*, April 2009, pp. 331–332. [Online]. Available: 10.1145/1530748.1530799

[88] L. Yang, Z. Zhang, W. Hou, B. Y. Zhao, and H. Zheng, "Papyrus: a software platform for distributed dynamic spectrum sharing using SDRs," *SIGCOMM Computer Communications Review*, vol. 41, pp. 31–37, January 2011. [Online]. Available: 10.1145/1925861.1925866

[89] P. Mackenzie, "Reconfigurable Software Radio Systems," Ph.D. dissertation, University of Dublin, Trinity College, 2004.

[90] Federal Communications Commision. [Online]. Available: http://www.fcc.gov/

[91] Office of Communications. [Online]. Available: http://www.ofcom.org.uk/

[92] K. Balachandran and J. H. Kang, "Neighbor Discovery With Dynamic Spectrum Access In Adhoc Networks," *Vehicular Technology Conference, 2006. VTC 2006-Spring. IEEE 63rd*, vol. 2, pp. 512–517, May 2006. [Online]. Available: 10.1109/VETECS.2006.1682877

[93] N. C. Theis, R. W. Thomas, and L. A. DaSilva, "Rendezvous for Cognitive Radios," *IEEE Transactions on Mobile Computing*, vol. 10, no. 2, pp. 216–227, February 2011. [Online]. Available: 10.1109/TMC.2010.60

[94] P. D. Sutton, K. E. Nolan, and L. E. Doyle, "Cyclostationary Signatures in Practical Cognitive Radio Applications," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 1, pp. 13–24, January 2008. [Online]. Available: 10.1109/JSAC.2008.080103

[95] P. D. Sutton, "Rendezvous and Coordination in OFDM-Based Dynamic Spectrum Access Networks," Ph.D. dissertation, University of Dublin, Trinity College, 2008.

[96] FFTW C FFT Libraries. [Online]. Available: http://www.fftw.org/

[97] B. Özgül, P. D. Sutton, and L. E. Doyle, "Bandwidth-adaptive MIMO OFDM for dynamic spectrum access networks," *Physical Communication*, vol. 2, no. 3, pp. 194–203, September 2009. [Online]. Available: 10.1016/j.phycom.2009.01.002

[98] "FCC Second Memorandum Opinion and Order, FCC Docket No 10-174," Federal Communications Commission, Tech. Rep., September 2010.

[99] "Implementing Geolocation: Summary of consultation responses and next steps," UK Office of Communications (Ofcom), Tech. Rep., September 2011.

[100] Boost C++ Libraries. [Online]. Available: http://www.boost.org/

[101] Rohde & Schwarz. FSVR Spectrum Analyser Data Sheet. [Online]. Available: http://www2.rohde-schwarz.com/file/FSVR_dat-sw_en.pdf

[102] CTVR - The Telecommunications Research Centre, "CTVR 2 Proposal," Science Foundation Ireland, Tech. Rep., 2011.

[103] T. Koponen, S. Shenker, H. Balakrishnan, N. Feamster, I. Ganichev, A. Ghodsi, P. B. Godfrey, N. McKeown, G. Parulkar, B. Raghavan, J. Rexford, S. Arianfar, and D. Kuptsov, "Architecting for innovation," *SIGCOMM Computer Communications Review*, vol. 41, pp. 24–36, July 2011. [Online]. Available: 10.1145/2002250.2002256

[104] C. Partridge, "Forty data communications research questions," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 5, pp. 24–35, September 2011. [Online]. Available: 10.1145/2043165.2043170

[105] I. F. Akyildiz, D. M. Gutierrez-Estevez, and E. C. Reyes, "The evolution to 4G cellular systems: LTE-Advanced," *Physical Communication*, vol. 3, no. 4, pp. 217–244, December 2010. [Online]. Available: 10.1016/j.phycom.2010.08.001

[106] M. Iwamura, K. Etemad, M.-H. Fong, R. Nory, and R. Love, "Carrier aggregation framework in 3GPP LTE-advanced [WiMAX/LTE Update]," *Communications Magazine, IEEE*, vol. 48, no. 8, pp. 60–67, August 2010. [Online]. Available: 10.1109/MCOM.2010.5534588