An Investigation on the use of SNR Distributions for the Optimisation of Coarse-Fine Spectrum Sensing for Cognitive Radio

Brendan Lawton

Department of Electrical and Electronic Engineering

University College Cork

A thesis presented to the National University of Ireland

Faculty of Engineering and Architecture

for the degree

Doctor of Philosophy

August 2013

Research Supervisor	:	Dr. Colin Murphy
Head of Department	:	Prof. Nabeel Riza

“A child of five would understand this, send someone to fetch a child of five.”

- Groucho Marx

Acknowledgements

First and foremost thanks must go to my supervisor Dr. Colin Murphy. Without his guidance, advice and support this work would not have been possible.

My family provided constant support, throughout not only the last four years but for the entirety of my education. They also kept me grounded by pointing out that I would not be a “real” doctor anyway.

To Claire, I would not have been able to get this work done without her constant support and endless enthusiasm for me finishing!

The Postgraduate Office must also be thanked, it was nice to know that I wasn’t going through this alone and that others were experiencing the same thing. Extra thanks must go to Donagh Horgan for a large number of technical chats, which helped enormously.

The Centre for Telecommunications Value-Chain Research in Trinity were absolutely fantastic throughout, always willing to help and offer advice, especially Dr. Keith Nolan and Dr. Paul Sutton. Without access to IRIS and the advice that they gave, this thesis would not be the same.

Finally, for when it all got too much, I would like to thank the UCC Squash club for providing stress relief.

Contents

List of Figures

Abstract

This thesis investigates the optimisation of Coarse-Fine (CF) spectrum sensing architectures under a distribution of SNRs for Dynamic Spectrum Access (DSA). Three different detector architectures are investigated: the Coarse-Sorting Fine Detector (CSFD), the Coarse-Deciding Fine Detector (CDFD) and the Hybrid Coarse-Fine Detector (HCFD).

To date, the majority of the work on coarse-fine spectrum sensing for cognitive radio has focused on a single value for the SNR. This approach overlooks the key advantage that CF sensing has to offer, namely that high powered signals can be easily detected without extra signal processing. By considering a range of SNR values, the detector can be optimised more effectively and greater performance gains realised.

This work considers the optimisation of CF spectrum sensing schemes where the security and performance are treated separately. Instead of optimising system performance at a single, constant, low SNR value, the system instead is optimised for the average operating conditions. The security is still provided such that at the low SNR values the safety specifications are met. By decoupling the security and performance, the system's average performance increases whilst maintaining the protection of licensed users from harmful interference.

The different architectures considered in this thesis are investigated in theory, simulation and physical implementation to provide a complete overview of the performance of each system. This thesis provides a method for estimating SNR distributions which is quick, accurate and relatively low cost. The CSFD is modelled and the characteristic equations are found for the CDFD scheme. The HCFD is introduced and optimisation schemes for all three architectures are proposed.

Finally, using the Implementing Radio In Software (IRIS) test-bed to confirm simulation results, CF spectrum sensing is shown to be significantly quicker than naive methods, whilst still meeting the required interference probability rates and not requiring substantial receiver complexity increases.

Associated Publications

- B. Lawton, C. Murphy, **Coarse-Fine Spectrun Sensing for Reduced Sensing Time**, 4^{th} International Conference on Signal Processing and Communications Systems (ICSPS), Brisbane Australia, December 13^{th} -15^{th}, 2010.

- B. Lawton, C. Murphy, **Minimizing the Coarse-Fine Spectrum Sensing Time for Cognitive Radios with Ideal Secondary Detectors Subject to Noise Uncertainty: An Analytical Approach**, 4^{th} International Conference on Cognitive Radio and Advanced Spectrum Management (COGART), Barcelona Spain, October 23^{rd} -26^{th}, 2011.

Acronyms

ADC Analogue-to-Digital Converter

AWGN Additive White Gaussian Noise

BEE2 Berkeley Emulation Engine 2

BER Bit-Error Rate

BPSK Binary Phase-Shift Keying

CAGR Compound Annual Growth Rate

CAV Covariance Absolute Value

CCDF Complementary Cumulative Distribution Function

CDF Cumulative Distribution Function

CDFD Coarse-Deciding Fine Detector

CF Coarse-Fine

CFD Cyclostationary Feature Detectors

CN Cognitive Network

CR Cognitive Radio

CREW Cognitive Radio Experimentation World

CSFD Coarse-Sorting Fine Detector

CTVR Centre for Telecommunications Value-Chain Research

DSA Dynamic Spectrum Access

DSM Dynamic Spectrum Management

DTV Digital Television

FCC Federal Communications Commission

FFT Fast Fourier Transform

GPS Global Positioning System

HCFD Hybrid Coarse-Fine Detector

ICDF Inverse Cumulative Distribution Function

IRIS Implementing Radio In Software

ISOM Incremental Self-Organising Map

KDE Kernel Density Estimation

LUT Look-Up Table

M-PSK M-ary Phase-Shift Keying

$MISE$ Mean Integrated Square Error

$MISC$ Mean Integrated Square Change

MSE Mean Square Error

OFDM Orthogonal Frequency-Division Multiplexing

QPSK Quadrature Phase-Shift Keying

QAM Quadrature Amplitude Modulation

ROC Receiver Operating Characteristic

SDR Software Defined Radio

SNR Signal-to-Noise Ratio

SOM Self-Organising Map

SU Secondary User

PDF Probability Distribution Function

P_{fa} Probability of False Alarm

PLL Phase Lock Loop

P_{md} Probability of Missed Detection

PU Primary User

USRP Universal Software Radio Peripheral

VCO Voltage Controlled Oscillator

WLAN Wireless Local Area Network

1

Introduction

1.1 Introduction

In recent years spectrum usage has increased dramatically. The proliferation of portable devices using mobile information services has caused spectrum scarcity issues. Cognitive Radio (CR) has been proposed as a possible solution to this problem. CR, as defined by Mitola [1], is an intelligent system capable of using contextual information to provide an improved service to the user. A CR should be capable of providing services based on

various conditions, such as location or spectrum occupancy. One of the most exciting possible applications of CR is Dynamic Spectrum Access (DSA).

DSA is the key technology that would allow CRs to solve the current spectrum scarcity problem. Using DSA a CR would be able to transmit in a licensed band of spectrum, provided the licensed Primary User (PU) is not interfered with. This ability should free, in many cases, a significant amount of spectrum for opportunistic use. It is estimated that current usage in licensed bands varies from 15% to 85% [2]. Therefore, a large amount of BW could be re-used with this method.

Termed the CR standard, IEEE 802.22 requires that, when the CR is sensing the channels, it should be able to detect signals with a SNR as low as -21 dB with a probability of missed detection (P_{md}) of, at most, 0.1 and a probability of false alarm (P_{fa}) of, at most, 0.1 [3]. This sensitivity ensures that the CR will not interfere with receivers at the edge of the primary network, where the primary user signal power is low. Being able to detect these very weak signals requires substantial signal processing. However, not all signals will require this level of receiver complexity.

For example, consider the situation illustrated in Fig. 1.1 where a CR scanning four bands as it is attempting to detect a spectrum opportunity, or free space, in one of the four bands. The ranges at which the CR is required to detect transmissions from each of the four PUs are shown for each of the bands. Only the fourth band is available for transmission. If the CR uses the full signal processing on all channels then it will detect the available band, or "spectrum hole". However, band one and band two contain higher powered signals and do not require such treatment. CF spectrum sensing can reduce the inefficiency of this method.

In a sorting based architecture, such as the Coarse-Sorting Fine Detector (CSFD) considered in Chapter 5, the receiver can gain some preliminary information about the spectrum

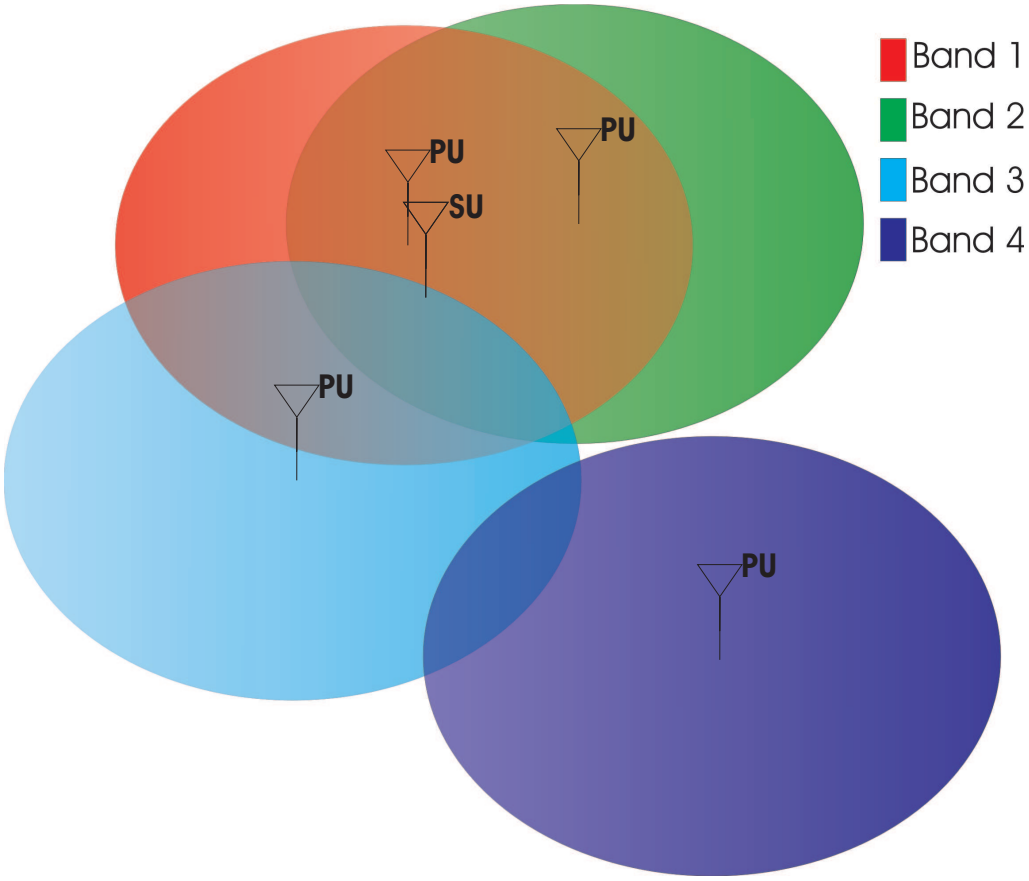

Figure 1.1: Typical Signal Environment with PUs at Different Distances to the CR

which helps it to choose the bands that are more likely to be unoccupied. When the signal sources are not equidistant from the receiver, such as in Fig. 1.1, the SNR will vary and significant sensitivity is not required to detected the presence of all the signals. The signal strengths at the CR vary from very high, Band 1, through medium strength, Band 2, to signals that are too weak to be detected and the bands are declared free, Band 4. If an estimate of the power in each band is found, then the bands can be ordered with respect to this estimate. As long as the noise power is constant across the bands, the bands without signals present will have a lower average power than the bands with signals present. This

estimate is found using an energy detector, as in Section 3.3, and, for the environment illustrated in Fig. 1.1, the output will resemble Fig. 1.2.

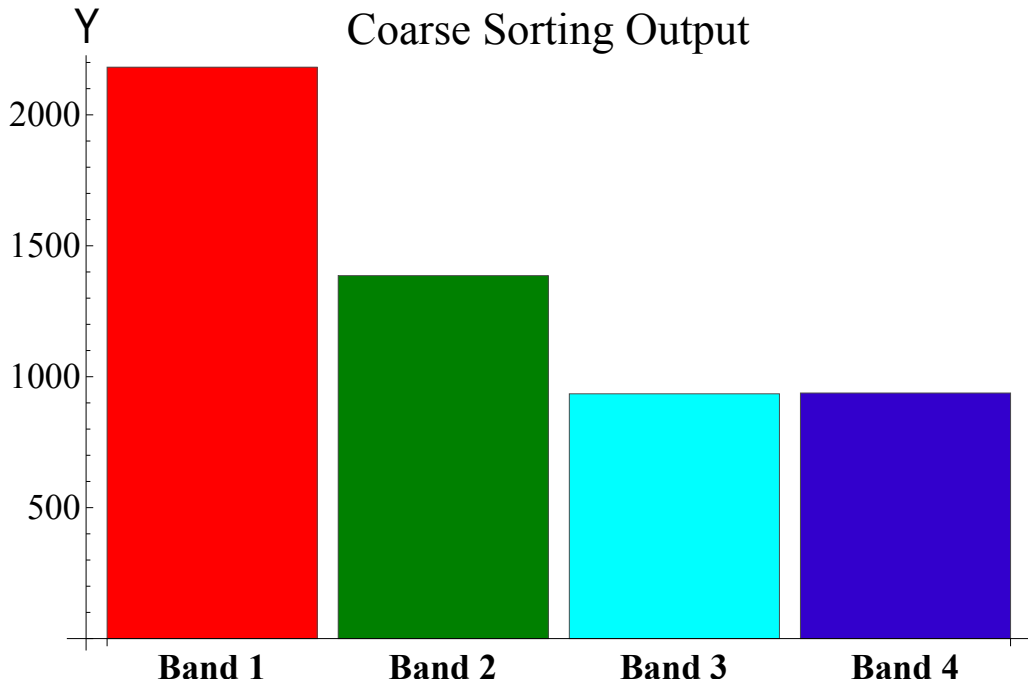

Figure 1.2: Energy Detector Output for Bands Under Sorting with 1000 Energy Detector Samples per band Used for Sorting

It is clear that either Band 3 or Band 4 is the most likely to be free, thus the detector would start its detection attempt in either channel three or channel four. This more informed selection of channel sensing order will reduce the number of detection attempts, on average, and, therefore, increase efficiency.

Another option for CF sensing is to exclude channels that are likely to be occupied. This can be thought of as deciding to remove the higher powered channels and is termed here a Coarse-Deciding Fine Detector (CDFD). This architecture is investigated in detail in Chapter 6. For example, if the detector sequentially checks for lower power signals and excludes any results where a signal exceeds a pre-defined threshold, then the overall effi-

ciency can be increased. This can be performed by taking increasing numbers of samples with an energy detector. The result of this method, for the environment in Fig. 1.1, is shown in Fig. 1.3 - Fig. 1.5.

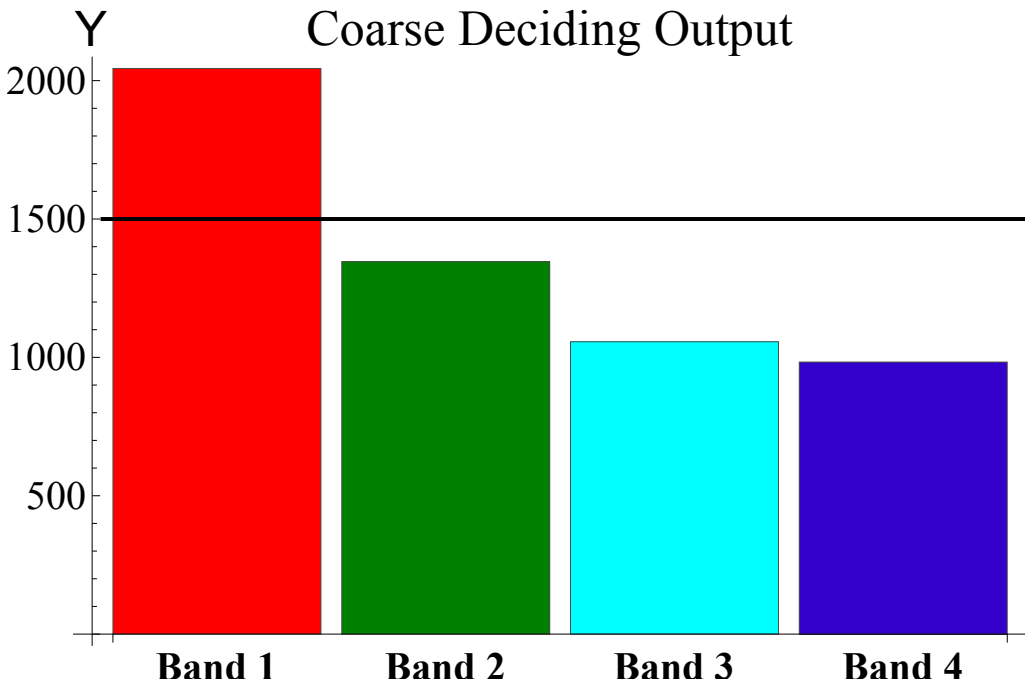

Figure 1.3: Energy Detector Output for Bands Under CDFD: First Detection Attempt

The first detection attempt uses 1000 samples, and the strongest signal, in Band 1, is detected and excluded. This is shown in Fig. 1.3.

The second detection attempt uses 10,000 samples and the next strongest signal, in Band 2, is detected and excluded. This is shown in Fig. 1.4.

The third detection attempt uses approximately 209,000 samples and this is sufficient to declare Band 4 free. This is shown in Fig. 1.4. Note that more than one channel can be excluded per attempt and also that each detection attempt does not have to exclude any channels.

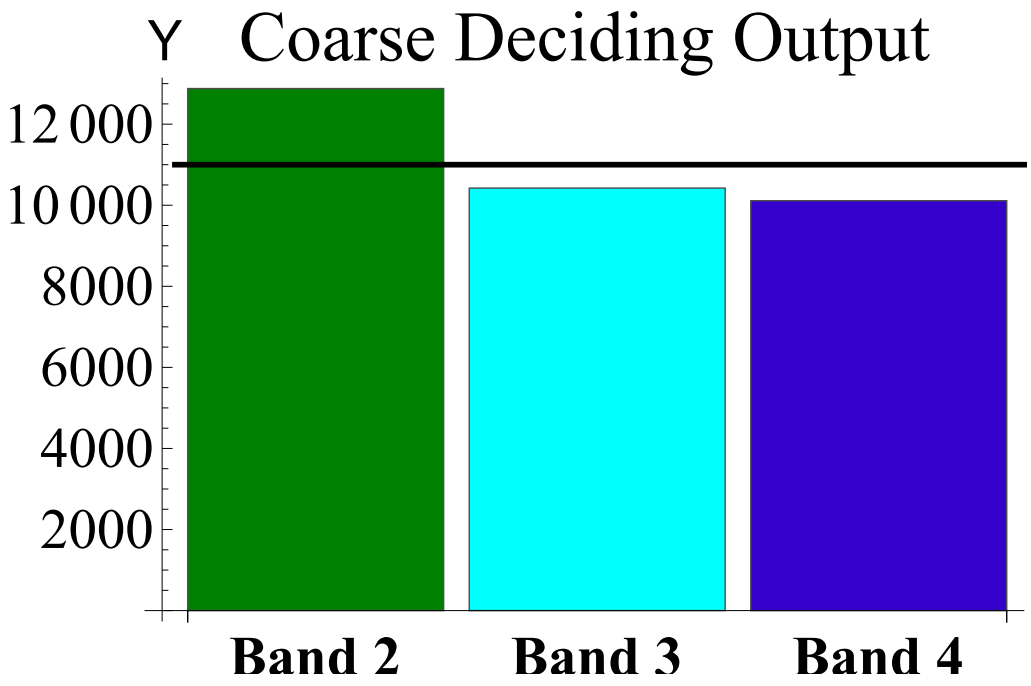

Figure 1.4: Energy Detector Output for Bands Under CDFD: Second Detection Attempt

CF sensing allows significant performance increases when a range of signal powers prevails. However, to the author's best knowledge, there has been no study published on the performance of CF sensing under these conditions and no method exists for predicting the optimum parameters for the sensing schemes, when a range of signal powers are present.

1.2 Objectives

The main objective of this thesis is to answer the following question:

> "How can a Cognitive Radio Coarse-Fine sensing scheme be optimised for the presence of a wide range of signal-to-noise-ratios, such that sensing time is reduced to a minimum, but interference probabilities remain unchanged, for Coarse-Fine schemes

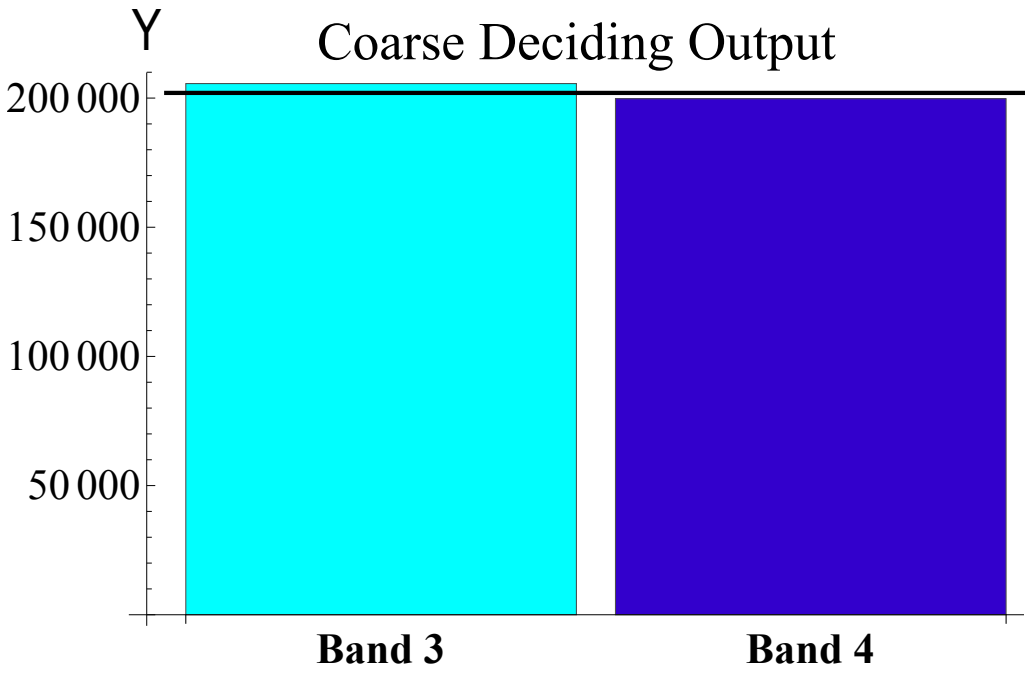

Figure 1.5: Energy Detector Output for Bands Under CDFD: Third Detection Attempt

based on sorting channels, deciding on channels or a combination of both?"

The optimisation of CF spectrum sensing will be investigated and optimisation schemes proposed for three CF architectures, namely CSFD, CDFD and a hybrid of the two, Hybrid Coarse-Fine Detector (HCFD).

1.3 Assumptions

Some assumptions have been made during the course of this work. Most are common assumptions when dealing with CR systems and are summarised here.

- The noise is assumed to be Additive White Gaussian Noise (AWGN) and a suf-

ficiently accurate estimate of the noise power is available. The signals are also assumed to have a Gaussian distribution.

- The noise power in the bands is approximately equal, or is known.
- The CR and PU can move in space but are assumed to be relatively stationary and that their locations do not change quickly relative to the sensing period. In addition, the PU transmissions are assumed to vary with periods significantly larger than the sensing time of the CR.
- The PU insists that the 10% false alarm and missed detection rates must still be met at -21dB, even if knowledge of the SNR distribution is exploited.
- The receiver has no additional knowledge of the conditions in the sensing environment, such as correlation in occupancy between adjacent bands and, thus, can only search randomly.
- This work is based on a single CR attempting to find free spectrum. It is assumed that there are no other CRs available with which the CR can cooperate to improve performance.
- For the CSFD it will be assumed that the signal powers in the bands are independent and identically distributed (i.i.d.). Whilst not strictly true, it reduces the computational complexity substantially, whilst not introducing any significant inaccuracy. This will be discussed further in Chapter 5.
- Finally, it has been assumed that the receiver front-end has multiple Phase Lock Loops (PLLs) or the receiver has a wide-band front-end, such that there is no significant time penalty when switching between channels.

1.4 Contributions

The contributions of this work are as follows:

1. A reliable method for generating an estimate of a SNR distribution using a CR attempting opportunistic access is shown.

2. A model of the CSFD is derived that is significantly quicker than Monte-Carlo simulations, whilst remaining sufficiently accurate over a wide range of practical conditions.

3. Optimisation equations and schemes for a CDFD are derived allowing a fast and accurate optimisation of the system.

4. A HCFD Architecture is introduced, using the CSFD and CDFD architectures, and is shown to outperform both methods.

5. It is shown that optimising these detectors using a uniform distribution can replace the need for in-line optimisation and SNR estimation, without reducing performance significantly, under the operating conditions investigated here.

6. Results were generated from an implementation of each of the receiver architectures on a test-bed that showed the architectures work in practice.

1.5 Outline

In Chapter 2, the current state of the art in CR is discussed. Starting with an overview of the field, various topics are discussed, focusing on spectrum sensing applications. The current options for CF spectrum sensing are reviewed and discussed. Also discussed are

the various test-beds currently in use and the CR system architectures used in each test-bed.

In Chapter 3, some basic theory is introduced. The basics of energy detector operation, including the issue of noise uncertainty, are investigated. The effects of fading channels and time varying channel occupancies are also considered. Markov Chain theory is also introduced and the relevant equations governing Markov Chains shown. Finally, the IRIS system and the Cognitive Radio Experimentation World (CREW) test-bed are examined and the IRIS architecture's structure shown.

In Chapter 4, SNR Probability Distribution Function (PDF) estimation is investigated. Various strategies for generating sample SNR PDFs are considered. In-line sensing is chosen as the most promising candidate and the advantages and disadvantages are shown. Testing and verification of the method is performed, both in simulation and on the IRIS system.

In Chapter 5, the CSFD architecture is considered. A new model of the CSFD is generated that matches Monte-Carlo simulations closely, whilst requiring significantly less (≈ 80 times) simulation time. By using order statistics to model the sorting operation and Markov Chains to model the effects of the sorting on the relevant probabilities of the fine detector, it is shown that the model predicts the CSFD performance accurately, even under fading and noise uncertainty conditions.

In Chapter 6, the CDFD architecture is considered. The characteristic equation of the CDFD is derived and three optimisation options investigated. It is shown that, by only allowing one false alarm rate to vary, performance close to the global maximum can be obtained, whilst reducing the complexity of the optimisation significantly.

In Chapter 7, the HCFD architecture is introduced. The HCFD is a combination of both techniques, CSFD and CDFD, and has better performance than either detector. It is shown

that the CSFD and the CDFD have close to optimal performance over a wide range of SNR distributions when optimised for a uniform distribution. Using this fact, the HCFD is not optimised directly, rather the parameters for the detectors optimised for the uniform distribution are chosen and the detector compared with the other architectures. This comparison is done both in simulation and using the IRIS test-bed for a practical implementation.

In Chapter 8, the work is concluded and future work proposed.

2

Literature Review

2.1 Cognitive Radio

In recent years there has been a substantial increase in the amount of data being sent wirelessly. For example, in 2010 the Federal Communications Commission (FCC) authorized nearly 12,000 wireless transmitters, almost four times that of 2000 [4]. This trend is set to continue with CISCO predicting that, “Global mobile data traffic will increase 18-fold between 2011 and 2016. Mobile data traffic will grow at a Compound Annual Growth

Rate (CAGR) of 78 percent from 2011 to 2016, reaching 10.8 Exabytes per month by 2016" [5].

This increase in demand has come at a time when the majority of available spectrum has been allocated. However, even when the allocation is high, the utilization typically remains significantly lower [6]. Spectrum utilization of 15% to 85% has been reported by some studies [7]. CR has been proposed as the solution to this spectral shortage problem [8].

The term CR was first coined by Joseph Mitola in 1999 [9]. Mitola describes a CR as an intelligent radio able to adapt to the needs of its user. The CR architecture is based on a Software Defined Radio (SDR), allowing greater flexibility, and a cognitive engine which adapts the radio to the situation [1].

Initially, a CR was defined as a radio which could:

- "Detect user communications needs as a function of use context"
- "Provide radio resources and wireless services most appropriate to those needs".

Some CR research has more recently focused more on the idea of Dynamic Spectrum Access (DSA) [10].

DSA occurs where unlicensed Secondary Users (SUs) are allowed to use spectrum owned by a licensed PU, provided the PU is not interfered with. DSA has become a central theme of CR research, and there are a number of technical challenges that need to be overcome before it can become a reality [11, 12]. The work in this thesis is concerned with fast reliable sensing for DSA.

2.2 Dynamic Spectrum Access

2.2.1 Hardware Requirements

Significant hardware challenges are present when designing a practical CR for implementation [2,13]. The IEEE 802.22 standard [3] requires that the CR be able to detect a signal at a signal to noise ratio (SNR) of -21dB, with a Probability of False Alarm (P_{fa}) and a Probability of Missed Detection (P_{md}) of less than or equal to 10% and that the entire sensing process be performed within two seconds. In addition, the fundamental principle underpinning CR is to enable the tailoring of signal characteristics to suit the situation. This flexibility comes at the cost of increased hardware complexity. The IEEE 802.11 af standard also has scope for opportunistic access but [14], though this has been envsioned as using a location awareness based scheme.

The radio front-end must be capable of signal detection over a wide range of frequencies but, at the same time, introduce little distortion. The IEEE 802.22 standard allows CR DSA in the bands between 41MHz and 910MHz. For efficient usage, it is proposed that the system employ a digital wideband receiver [15]. If an analogue filter with a narrowband frontend was used, then the system would need to change the frequency for each new band being scanned. A PLL would have to be tuned to the new frequency each time and the settling time required for the PLL would most likely increase the required sensing time significantly [16].

For the wideband receiver it has been shown that Analogue-to-Digital Converters (ADCs) play an important role in determining receiver efficiency [17]. Hardware imperfections and quantization noise reduce the efficiency of the ADC. The dynamic range required is quite large, the SNR can be as low as -21dB in some bands but it is possible that it could be as high as 20dB in others, where the signal source is nearby and strong. In

addition, the wideband architecture will, in practice, reduce the effective SNR further before sampling.

Other hardware implementation issues include non-linear effects from the Voltage Controlled Oscillator (VCO) used by the PLL [18]. Typically, harmonics are generated by the VCO, at odd multiples of the fundamental frequency, that could cause distortion capable of compromising the detection algorithm. Finally, all of these hardware problems need to be solved by a portable device having low power requirements.

2.2.2 Self Organising Networks

One advantage that a network of CRs, or a Cognitive Network (CN), has over traditional systems is the ability to self-organise [19]. A CN differs from a CR in the scope of the parameters that can be changed. A CR is mainly concerned with the physical layer and the link layer of the OSI seven layer model. A CN can optimise over the entire operating conditions. This is illustrated in Fig. 2.1, where the scopes of the two concepts are shown on the OSI seven layer model.

Self-organisation can occur when a system is allocating frequencies for individual radio nodes to use. Self-Organising Maps (SOMs) [20] can be used as a method of Dynamic Spectrum Management (DSM). By creating a SOM for the nodes using observations of their local signal environment, it is possible to reduce the probability of interference with PUs whilst also allowing the network to communicate efficiently. In addition, the system is computationally simple, based on the Hebbian learning [21] (associated learning) rule. It has been shown that SOMs can improve DSM, though the specific amount will depend on the network conditions [20]. Note that the self-organising networks are used for a distributed network where there is no central controller.

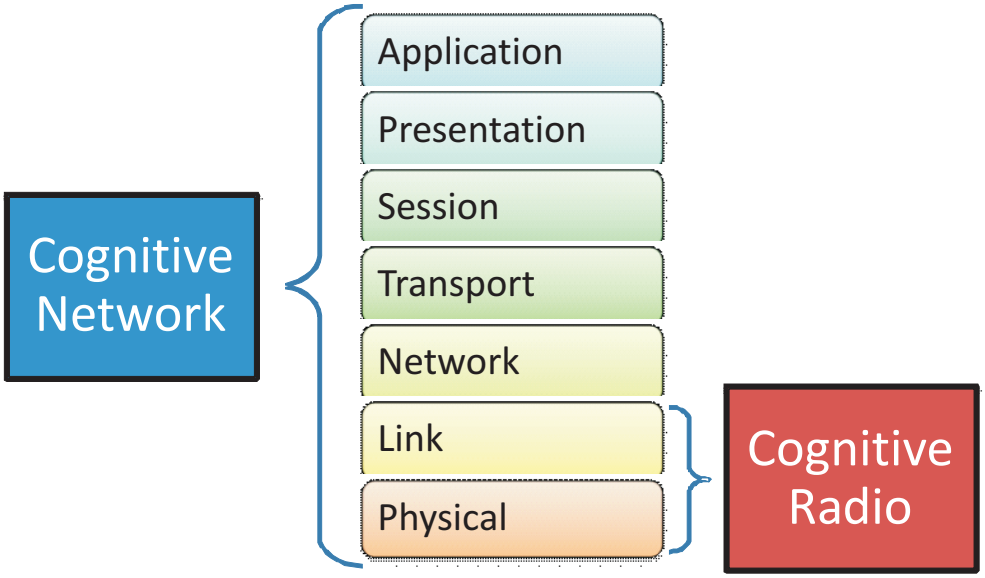

Figure 2.1: OSI 7 Layer Model and the Scope of CN and CR

An improvement on SOMs for CR applications is the Incremental Self-Organising Map (ISOM) [22]. ISOM has an intelligent weighting system that allows the system to learn, starting with a total lack of information, which is the expected initial condition of a CR. In addition, ISOMs allow the weighting system to be changed to facilitate shorter or longer learning periods, depending on the prevailing radio conditions and user requirements.

One threat to self-organising networks is malicious users who seek to break the rules of spectrum sharing [23]. A CR could lie about its received signal environment to other users in the network, thus keeping free spectrum for itself by declaring it occupied. Alternatively, another user may attempt to create interference for the PU by declaring an occupied channel free. Other possibilities for attack stem from the imitation of a primary user to prevent other CRs from attempting to share spectrum [24]. A CN would have to be robust to such attacks to be commercially viable.

When a CR attempts to access free spectrum, there are two main options for the allocation

of frequency bands. Firstly, there is the case where the PU is aware of interest in allocating its band and holds an auction to decide which CR is given access [25,26]. Each CR can bid for the spectrum and the auction winner can subsequently use the spectrum for a set period of time. This allows the PU to profit from the opening of the spectrum to a CR which, it is hoped, would help incentivise PUs to release bands to CRs for DSA applications. One drawback of this system is the requirement that the PU oversees the spectrum before an auction can be held. Thus, there is no free method for including this option in bands occupied by legacy PU systems [27].

Another option is where the secondary users decide to allocate the bands fairly between all users of the network, depending on each CRs individual need [28]. A centralised node can decide on the allocation of the bands such that each CR obtains fair access. If one node requires more bandwidth than the other nodes and there is spectrum available, it will allow the node to transmit with greater bandwidth. One challenge with this method is security. If a node lies about its requirements, either to damage the network or to ensure that its lower requirements are fully met, then the overall system performance can be degraded.

2.3 Spectrum Awareness for DSA

To enable DSA, spectrum awareness remains the key issue. If the CR does not have knowledge of the prevailing radio environment, it cannot guarantee that its transmissions will not interfere with a PU which is, clearly, unacceptable. In general, the CR has to be more sensitive than the PU if it is to ensure that little or no interference occurs. For example, consider the situation illustrated in Fig. 2.2. The CR has an obstruction between it and the transmitting PU and receives a low power signal. The receiving PU does not

have any obstruction and receives a relatively high powered signal. If the CR declares the channel free and transmits, then it might cause interference to the PUs in the band. This is termed the Hidden Node Problem.

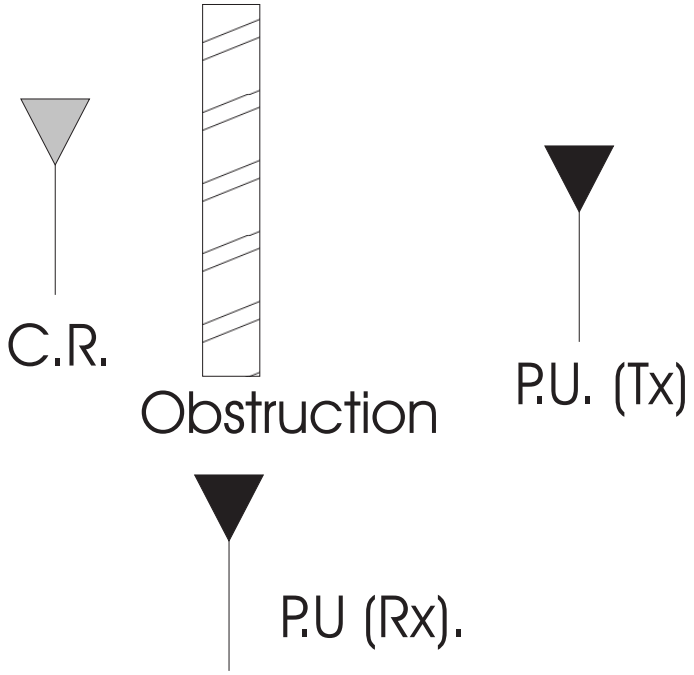

Figure 2.2: Hidden Node Problem for Cognitive Radios

Another situation is shown in Fig. 2.3. The distance over which a receiving PU is able to detect a transmitting PU is shown by the area centred on the transmitting PU. The CR is outside this region. However, if the CR transmits, then it will interfere with the PU at the edge of the transmitting PU's range. Therefore, an extra exclusion range is required.

In Cabric 2004 [2], the authors present a review of the requirements for spectrum sensing in a CR system. The authors also summarize and compare possible sensing algorithms. They show the inherent advantages and drawbacks of each method; Matched filters, energy detectors and Cyclostationary Feature Detectors (CFD). Also shown is the potential

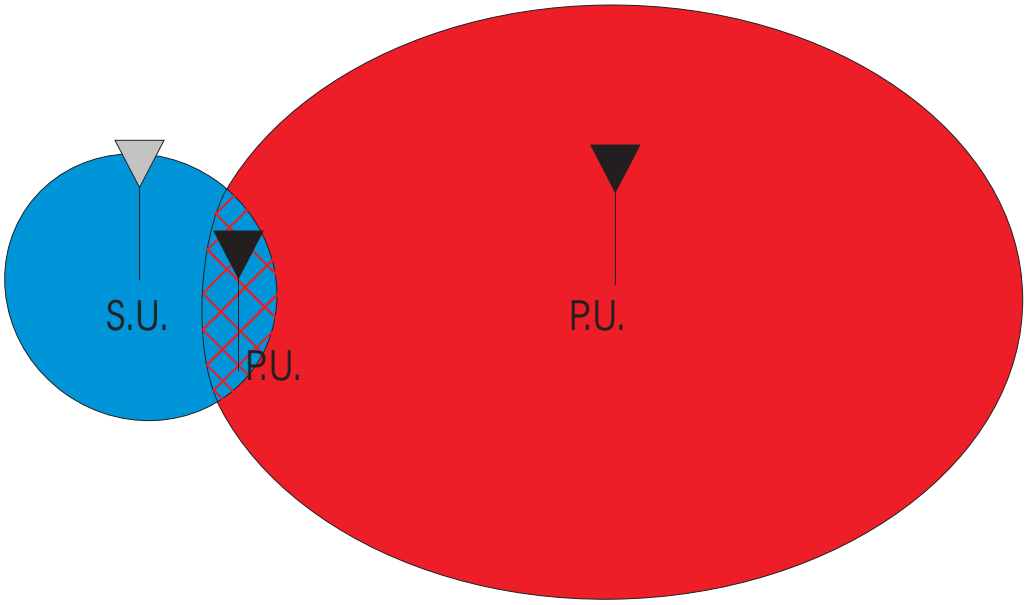

Figure 2.3: Requirement for Extra Sensing for Cognitive Radios

of co-operative spectrum sensing to increase reliability whilst also reducing sensing time. The authors generate a signal environment, in simulation, that is used to show the gain in performance that cooperation between CRs can provide. The co-operative system is compared to several individual radios, which are not cooperating, and a significant performance gain is shown. These tests were also performed without any optimisation of the voting rules for the co-operative network and thus, as discussed in 2.3.2, the performance of the system would increase dramatically with optimisation.

To solve the challenge of DSA for CRs, three possible solutions are proposed.

1. Location Awareness

2. Co-operative Spectrum Sensing

3. Single Radio Sensing

2.3.1 Location Awareness

Due to the inherent problems associated with spectrum sensing, an alternative method is to avoid it entirely. With location based spectrum awareness, the CR uses Look-Up Tables (LUTs), based on its location, to determine which bands are available for use, and at what transmit power. These LUTs are often known as geolocation databases [29].

In Fig. 2.4 an example of a location awareness scheme is shown. There are seven locations, each with two possible channels that might be free. If the CR is in location one or location five, then it can use band one. If it is in location four or location six, then it can use band two. If it is any other location, it can use either of the bands.

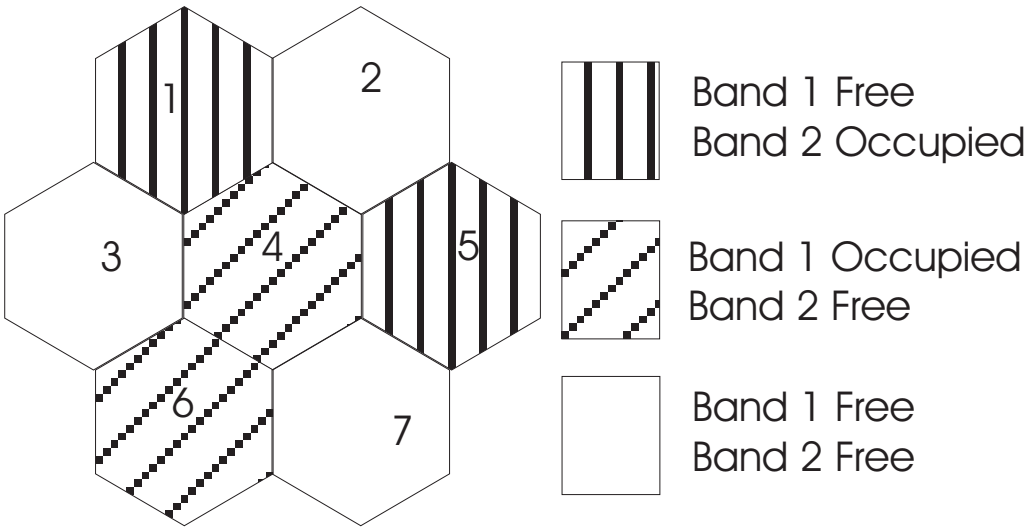

Figure 2.4: Example Location Awareness Scheme

However, a difficulty for geolocation based systems is presented by quickly changing bands. If the occupancies of the bands change rapidly, for example in the Wireless Local Area Network (WLAN) band, then the database has to be updated frequently. This also requires the PU to inform the databases of its change, a problem for legacy technology as

new components would be required on all transmitters.

Despite these issues, for slowly changing bands, such as Digital Television (DTV) bands, geolocation databases are an effective alternative to spectrum sensing [30]. Typically, DTV bands are very slow to change occupancy or transmitter location, thus the database would not require frequent updating. DTV bands generally cover large spatial ranges, thus the number of PUs required to identify themselves is lower. Also, since DTV channel allocations and transmitter locations are generally static, it is significantly easier for them to communicate with the database if changes to the database are necessary.

One final concern for geolocation based systems is the requirement for location information within the CR. If the CR requires a Global Positioning System (GPS) detector to determine its location accurately, then the complexity savings accruing from avoiding a spectrum sensor are largely lost when implementing the GPS detector. In addition, in urban environments, GPS signals are typically much less reliable than in rural areas [31].

2.3.2 Co-Operative Spectrum Sensing

Co-operative spectrum sensing occurs when multiple CRs attempt to sense the spectrum and aid each other in generating the result. This can be seen in Fig. 2.5. Even though two of the five CRs are obstructed, the system should be able to detect the presence of the PU and avoid interference.

For a network of co-operating CR nodes, there are a number of methods for determining the occupancy of the channels. A simple method is to give each node a separate channel to scan. This allows multiple bands to be scanned quickly but does not give the same advantages of other methods, such as the improvement in fading robustness shown in [32]. Other methods require the use of a fusion centre. A fusion centre is a CR where all the

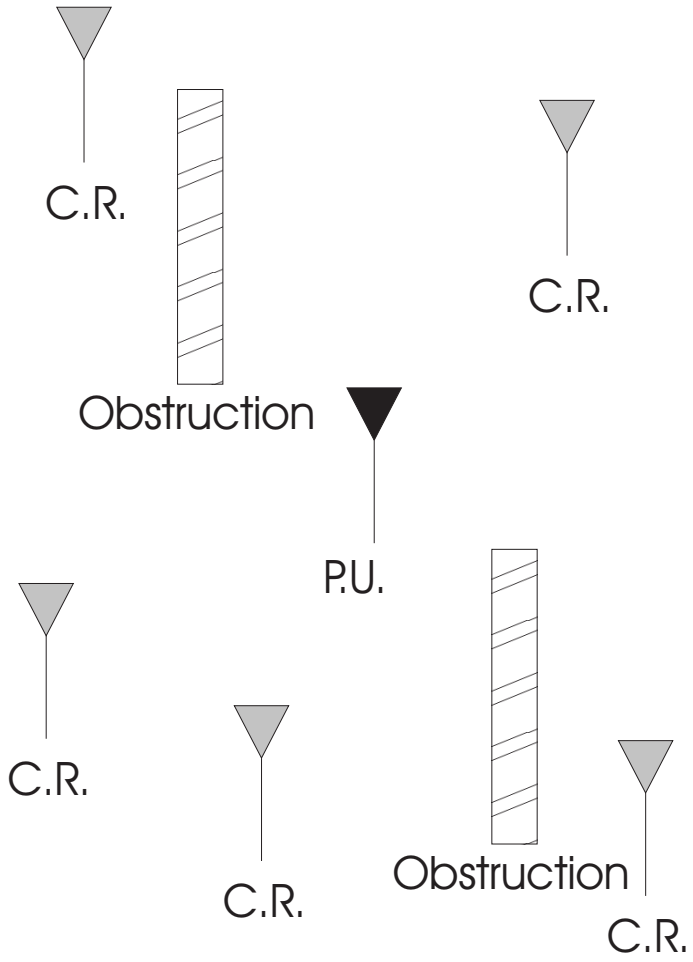

Figure 2.5: An Example of Co-Operative Spectrum Sensing

results are collated to make the decision [33]. This central node then communicates the result and, therefore, the occupancy to all the other radios.

The simplest sensing scheme for co-operative sensing is called the voting rule, or the counting rule [34]. Each node performs its detection separately and sends its decision to the fusion centre, for example either “1”, if it decides there is a signal present, or “0” if it decides there is no signal present. Then, the fusion centre counts the number of nodes that decide a signal is present and compares this to a threshold. If it exceeds the threshold,

then the system declares the band occupied. It has been shown that the "half voting rule", where the threshold is set to be equal to half the number of co-operating nodes, is optimum for this method [35].

Instead of sending just the binary decision, sometimes called hard fusion, if the nodes send the numerical result of the detection, then improved performance can be attained [36]. This method of co-operative sensing is called soft fusion. The fusion centre then uses this information to create a likelihood ratio test that satisfies the conditions for DSA. It has been shown that soft fusion significantly outperforms hard fusion.

However, soft fusion requires sending more information than hard fusion. The analysis performed in Visotsky [36] was for infinite precision soft fusion. For a practical system, the number of bits transmitted must be limited. To the author's best knowledge, there has been no agreement on the required number of bits but it has been suggested that it is not significantly greater than that of hard fusion [37].

A significant advantage of co-operative sensing over single radio sensing is found when fading is considered. It has been shown that co-operative sensing reduces the effect of fading due to the averaging effect on the fading across the nodes. This sensing diversity gain is present provided the radios are sufficiently far apart to be subjected to independent fading [32]. For example, twenty radios in Rayleigh fading at an SNR of -21dB require approximately 1000 samples each, whereas a single radio would require approximately 4.8M samples, to guarantee acceptable performance [32].

Co-operative spectrum sensing has been shown to mitigate noise uncertainty [38]. By the use of a double threshold energy detector at each node, the effects of noise uncertainty can be reduced. Each of the individual nodes uses a double threshold energy detector which has three possible results, namely Signal Present, No Signal Present and Uncertain (where the detector decides that it does not know the occupancy of the channel). If the

width of the uncertainty region matches the noise uncertainty, it is possible to reduce the uncertainty in the results forwarded to the fusion centre, in a hard fusion scheme. In a soft fusion scheme the gain would be analogous to the averaging across the nodes in the case of fading. Once again, this requires that the noise uncertainty be independent across the nodes [39].

Co-operative spectrum sensing, of course, requires multiple radios to be present, and that all the radios are willing to co-operate. In the case of a single radio, or where radios are too far away from one another to co-operate, co-operative spectrum sensing cannot be used and the CR must determine the state of the spectrum independently. In addition, there may be a redundancy of information from some of the nodes, a condition that, to date, has not been explored in the literature for many of the algorithms. In this work, the fundamental sensing case will be investigated, i.e. a single radio attempting to find the first available spectrum opportunity.

2.4 Single Radio Sensing

For a single CR attempting to find a spectrum opportunity there are a number of detection options for the sensing scheme [40], namely:

1. Matched Filters
2. Energy Detectors
3. Cyclostationary Feature Detectors (CFDs)
4. Other Detector Architectures

2.4.1 Matched Filtering

Matched filters are the most efficient method of detecting the presence of a signal [2]. The required number of samples for a matched filter grows at rate of $O(1/SNR)$ samples at low SNR [41] for static P_{fa} and P_{md}. However, the receiver requires perfect knowledge of the PU signal characteristics, such as bandwidth, modulation type, operating frequency and any other relevent transmission characteristics. In practice, it is highly unlikely that this information will be available to the system. In addition, the CR would require a separate receiver for all signal types, which is highly impractical.

2.4.2 Energy Detectors

Energy detectors are a commonly studied detector type for CR applications [18, 32, 34–37, 42–53]. They are simple to implement, requiring little computation and no knowledge of the signal characteristics, other than an estimate of the noise power. Energy detectors require more samples for detection than matched filtering, with the required number of samples growing at rate of $O(1/SNR^2)$ samples at low values of SNR [40], for static P_{fa} and P_{md}. This makes them ideal for co-operative sensing, where the expected number of samples required is low. Also, most CF algorithms use energy detectors as the coarse detector. This will be discussed further in Section 2.5.

An energy detector operates by estimating the energy in a band and comparing it to a threshold, λ [54]. The output of the energy detector, Y, is:

$$Y = \sum_{n=1}^{N} \mid x(n) \mid^2, \tag{2.1}$$

where $x(n)$ is the n^{th} sample and N denotes the total number of samples taken for the test. The required number of samples for an energy detector, N, is given by (2.2). The

energy detector is addressed in detail in Chapter 3 and is subsequently used extensively throughout this thesis.

A fundamental limit of energy detectors, which may impact on their usability in CR applications, is that of noise uncertainty and SNR walls (an SNR wall is an SNR value below which it is impossible to guarantee performance). Work has been done to investigate the effects of inaccurate noise power estimation for energy detectors [42]. If the estimate of the noise power is wrong by a factor of ρ ($\rho \neq 0$) then, where previously the equation for the required number of samples was:

$$N = \frac{2[Q^{-1}(P_{fa}) - Q^{-1}(1 - P_{md})(1 + SNR)]^2}{SNR^2}, \tag{2.2}$$

and Q^{-1} is the inverse Q function, the new equation becomes:

$$N = \frac{2[Q^{-1}(P_{fa}) - Q^{-1}(1 - P_{md})]^2}{(SNR - (\rho - \frac{1}{\rho}))^2}. \tag{2.3}$$

As the SNR approaches $(\rho - \frac{1}{\rho})$, N will increase towards infinity, creating what is termed an SNR wall. This sets a lower bound on the SNR that an energy detector can use to reliably detect signals. With a noise uncertainty of 0.1 dB, it has been shown that SNRs of less than -14dB cannot be reliably detected. To detect a signal with an SNR of -21dB a noise uncertainty, ρ, of $\leq$ 0.017dB is required. This presents a significant challenge to DSA under the IEEE802.22 standard. Noise uncertainty is addressed in more detail in Chapter 3. However this is only for an single radio scheme using an energy detector. A single radio using a cyclostationary feature detector does not suffer from noise uncertainty. In addition, a co-operative scheme using an energy detector can reduce the impact of noise uncertainty dramatically, though the exact amount has not been quantified. An energy detector in the presence of fading suffers significantly degraded performance

[55]. If the signal previously had an SNR of γ dB, then the faded SNR becomes a distribution with a mean value of γ. By averaging over this new SNR distribution, it is possible to generate an expression for the new P_{md} of the system and, thus, the number of samples required. For a signal under Rayleigh fading, at an average SNR of -21dB, the required number of samples for a P_{fa} and P_{md} of 10% is 4.8M, compared to 209k samples when fading is not present. Fading is addressed in more detail in Chapter 3.

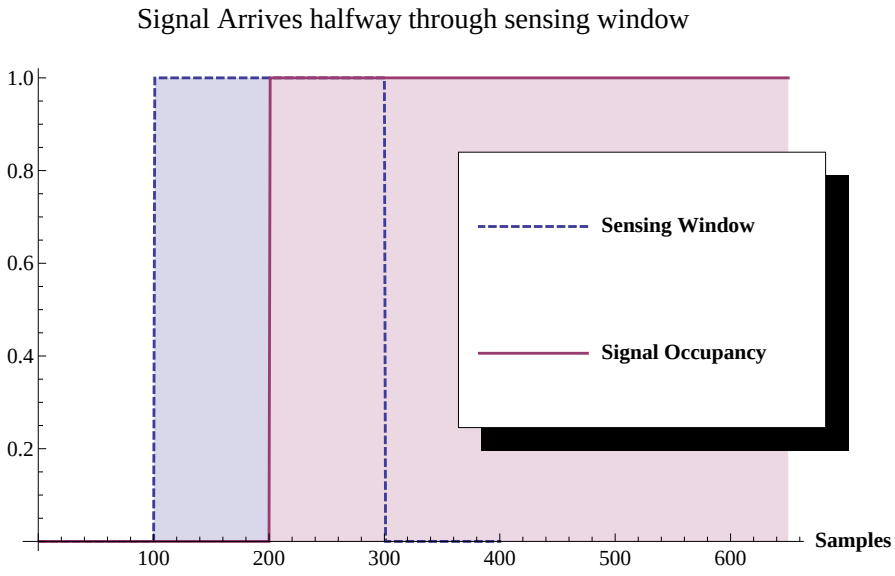

Figure 2.6: A Signal Appears During the Sensing Interval

These equations assume a static environment, where the signal is either present or not present and the occupancy does not change during the sensing interval. If this is not the case, then a probability based scheme would have to be introduced, as in Ma 2008 [43]. Here the case where a PU begins transmitting during the sensing interval is investigated. The situation is illustrated in Fig 2.6 where a PU begins transmitting approximately mid

way through the sensing attempt. If $x(n)$ denotes the n^{th} sample of the received signal then, if the PU begins transmitting at the m^{th} sample, $x(n)$ is given by:

$$x(n) = \begin{cases} v(n), & 1 \leq n \leq m \\ s(n) + v(n), & m < n \leq N \end{cases}, \tag{2.4}$$

where $v(n)$ is the n^{th} noise sample and $s(n)$ is the n^{th} signal sample.

If the distribution of m is known it is possible to generate a probability based scheme that weights the samples based on the probability that a signal will appear before that sample. This scheme has been shown to perform better than a conventional energy detector scheme for signals that can appear within the sensing attempt, though an analysis under different conditions, such as when the signal disappears after a certain period, was not investigated. This architecture will be proposed later in Chapter 5 and Chapter 6 as a possible solution when the time varying channel significantly degrades performance.

One important question is the improvement of the system under the condition that the signal source stops transmitting during a sensing attempt. The weighting would then improve the probability of declaring the channel free but, to the author's best knowledge, the characteristics of this improvement have not yet been quantified in the literature. However, this is not as important a parameter since the main concern of CR detection schemes is avoiding interfering with PUs.

In Gahasemi 2007 [45], a formulation for the time required for sensing is derived, for an energy detector based scheme. The sensing time is set as the time taken to find the first available channel of N_{ch} channels. The average search time, $\bar{T}_{search}$, is given by:

$$\bar{T}_{search} = \frac{T_s}{P_a}, \tag{2.5}$$

where T_s is the search time per band and P_a is the average probability of a channel being declared unoccupied and used for transmission. P_a is a function of the detector constraints and the occupancy profile of the channels.

Both of the systems in [43, 45] assume an exponential distribution for the channel occupancy duration. This is based on experimental results generated from separate studies, in [56, 57]. The occupancy probability, P_{occ}, after t seconds is:

$$P_{occ}(t) = \lambda e^{-\lambda t}, \tag{2.6}$$

where λ is the mean time for switching between states. For quickly changing channels, such as Wi-Fi, λ could be very small, of the order of milliseconds or seconds. For more slowly changing channels, such as DTV channels, modelling the change in occupancy is less necessary. Under normal circumstances, they can be approximated as static channels when performing individual sensing attempts as $\bar{T}_{search} << \lambda$. This is investigated further in Chapter 3.

Traditionally, energy detector performance has been specified by setting the P_{fa} and then minimising the P_{md}, such that the sensing requirements are met. However, the detector could be designed such that P_{md} is set and then P_{fa} is minimised [46]. This, it is argued, reduces the interference of the CR on the PU, though the CR requires knowledge of the PU's SNR. However, if P_{fa} and P_{md} are set at the threshold values for the lowest required SNR, as in conventional schemes, then the P_{md} will be lower for all other values of SNR. This results in a smaller probability of interfering with a PU than the case where P_{md} is set to the threshold value.

Although setting P_{md} for an individual SNR value each time will realise a speed gain, requiring knowledge of the SNR renders it less practical. In Section 2.7, SNR estimation

will be reviewed.

2.4.3 Cyclostationary Feature Detectors

CFDs rely on the inherent periodicity of man-made signals to detect these signals in the presence of noise [58]. A cyclic feature will only appear where one of the statistical features of a signal, such as the mean or the autocorrelation, is periodic. CFDs have several advantages over energy detectors and have been studied quite extensively for use in CR applications [59–69].

One major advantage of a CFD is that different signal types have different cyclic features. It is possible to train a CFD to recognise the difference between modulation types [62, 63, 68]. This allows a CR to distinguish between different users in a network and identify the PU. However, to the author's best knowledge, the ability of a CFD to perform the identification has not been published for low SNR values and this may reduce the performance dramatically.

Cyclic features can be embedded intentionally into signals to aid with system control [59–61], even in the presence of frequency selective fading. This would allow a system to re-organise itself quickly in the event of a PU returning to the band, without requiring a control channel.

CFDs do not require an estimate of the noise power, thus noise uncertainty is not a problem [69]. This is a significant advantage over energy detectors. In addition, the performance of CFDs can be improved upon by various means. Implementation time can be reduced by performing a 1^{st} order cyclic test, though this requires knowledge of the signal characteristics for efficient operation [65]. If the cyclic frequency is low, then decimation of the frequency to allow more samples to be taken can improve detector performance

significantly [66].

A subset of CFDs is the multitaper method. The multitaper method has been shown to outperform cyclostationary feature detection, whilst still providing all of the features that make a CFD so attractive for CR applications [67, 70], especially when filter banks are used [71, 72].

The main disadvantage of CFDs is the computational complexity of the detector. Unless information is known about the signal, a two dimensional search space is required. The search is required in both the cyclic frequency and the time delay dimensions. The algorithm itself is also more complex than an energy detector, though it can be accelerated by Fast Fourier Transforms (FFTs) [66]. Thus, the system requires significant processing power to detect a signal. While CFDs are not directly studied in this thesis, it is a possible option for the fine detector used in Chapter 5 and the model developed can be used with a CFD fine detector, if required.

2.4.4 Other Detector Architectures

Several other detector architectures have been proposed to meet the sensing challenges of IEEE 802.22. One proposal for the sensing scheme is to have a pro-active scheme that probes the band in question with the aim of increasing detection rates [73]. If the PU has an active power control scheme then increasing the interference in that band will increase the transmit power, thus making detection easier. However, since spectrum sensing is required to enable the CR to avoid interference, there is a question of whether PUs will be open to a scheme that requires generating interference to avoid causing interference.

Another option to reduce sensing time is to increase the number of antennae on the system [16, 44, 74]. However, if the total number of antennae is increased on a single radio there

will be a correlation between the results such that performance gain is sub-linear in the number of antennae. Indeed, it seems very unlikely that 100 antennae would need to be implemented on a single radio [44]. A better performance gain accrueing with multiple antennae is realised when statistical covariances are used [75]. However, once again, the requirement that the antennae be sufficiently separate spatially such that the noise is uncorrelated, is difficult to achieve in practice, especially on a handheld device. It is worth noting that the analysis presented in [75] can be used for co-operative sensing where the radios are connected through a broadband control channel. It is highly likely that this would reduce the problem of correlation significantly.

Wavelet transformations can be used instead of traditional transformations to improve upon the performance of the detector. A Wavelet transformation allows edge detection that can be more efficient at finding the occupied and unoccupied bands. The detector identifies all the locations at which the PSD changes significantly, usually denoting the edge of a band. This information allows the CR to select the bands that have lower PSD and are, therefore, more likely to be unoccupied [47]. This method is more appropriate for selecting promising candidate bands for another detector type, such as a CFD and, therefore, wavelets will be discussed again under the C-F detector architectures in Section 2.5.

Compressive Sensing allows CRs to detect signals when sampled below the Nyquist rate [76, 77]. However, it is quite computationally intensive and requires significant signal processing. When optimised for a single signal, multi-resolution Bayesian Compressive Sensing can be performed more quickly and with fewer samples than normal Compressive Schemes [78] but, again, the computational overhead is large.

Detector performance can also be improved by changing the search scheme used [79, 80]. Consider Fig. 2.7, where the channels are correlated in occupancy with a correlation of

Δ. If a channel is occupied then the probability that the next channel is also occupied is Δ. Thus, if a channel is occupied, the adjacent channels are more likely to be occupied, for $\Delta > 0.5$. Traditional methods are a random search scheme, where the channel is selected at random or a serial search, where the channels are selected sequentially. For high correlation values (e.g. in the range $\Delta \geq 0.9$) a scheme called the n-step serial search outperforms both methods. Instead of sequentially selecting the next band, the scheme skips the next n bands. This reduces the probability that the channel scanned is in a similar state to that of the previous channel. The scheme performed better than the random search for a channel correlation of $\Delta \geq 0.9$ where the occupancy was 70%.

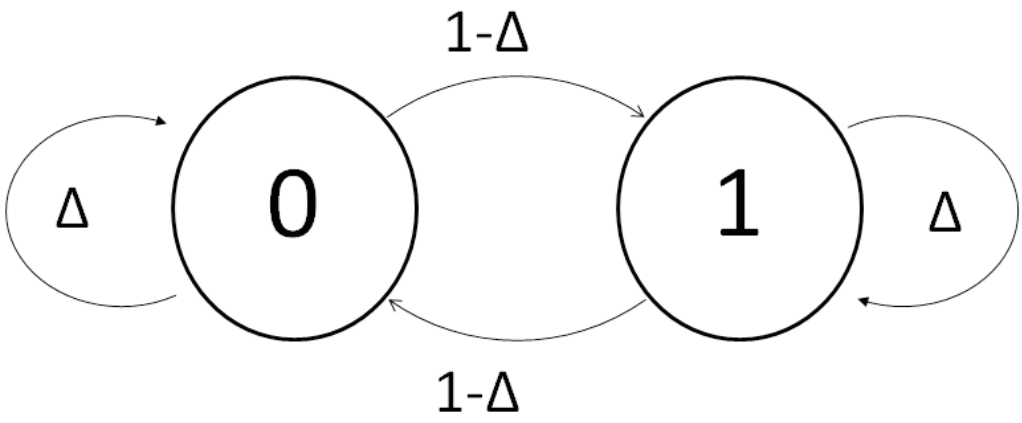

Figure 2.7: Markov Model of Correlated Channels

It is argued in [79,80] that, without additional information, this is the most efficient search scheme that can be implemented. However, CF sensing allows additional information to be generated that increases the efficiency of the search.

2.5 Coarse-Fine Detectors

An obvious compromise between the energy detector's speed and the CFD's ability to detect weak signals is to use the energy detector first to find the high SNR channels and then the CFD to decide between low SNR channels and unoccupied channels. This type of detector is called a CF detector or a multi-resolution detector. The more accurate of the detectors, usually called the fine detector, can be of any suitable architecture (eg. a CFD or an energy detector), provided the interference constraints are satisfied. The architectures can be divided into two types:

1. The coarse detector sorts the channels for the fine detector [18, 47–50].

2. The coarse detector decides on the occupancies of some of the channels and then the rest are scanned by the fine detector. [16, 51, 52, 81–84].

A sorting based CF detector accelerates the radio's attempt to find the first available channel. The sorting of the channels allows the fine detector to make better decisions regarding which channels to scan first. It does not, however, reduce the overall time to find the occupancies of all the bands. Thus, it is more applicable to situations where finding a channel to transmit is the main goal.

A deciding based CF detector reduces the time to find the occupancies of all the bands. The bands containing high SNR signals are declared occupied by the coarse detector and the fine detector is only run on bands about which the coarse detector is uncertain. Generally, this reduces the time to find the first band. It is generally more applicable for situations where the CR requires complete spectrum knowledge.

Both CF sensing types are investigated in this work. The sorting based scheme (denoted the CSFD) is investigated in Chapter 5, the deciding based scheme (denoted the CDFD)

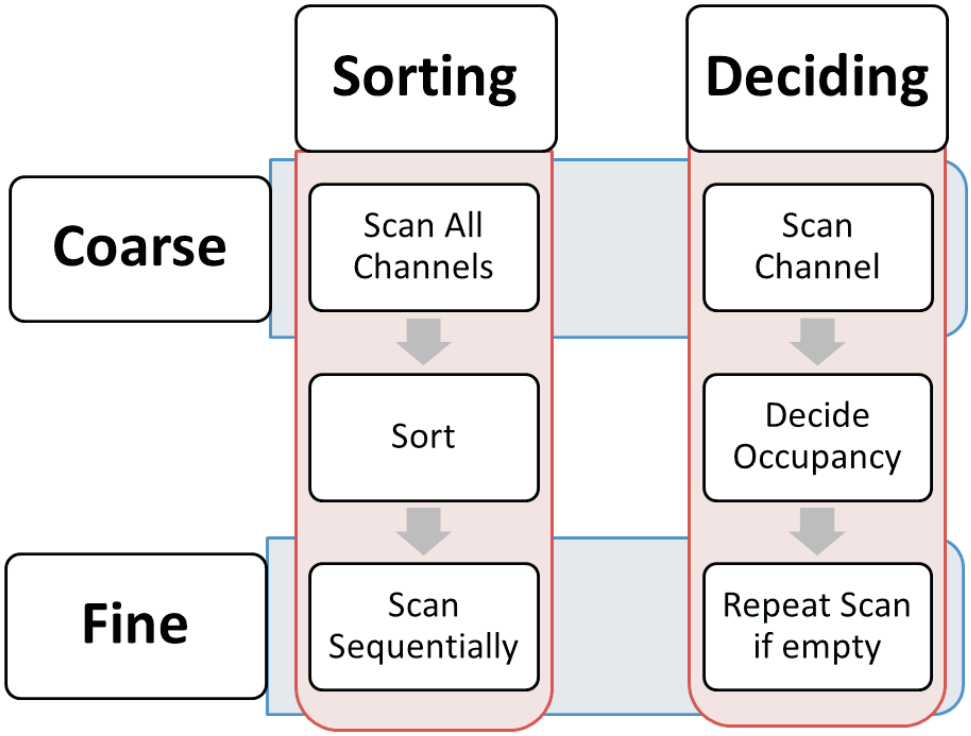

Figure 2.8: Coarse-Fine Architectures

is investigated in Chapter 6 and a hybrid of both (denoted the HCFD) is investigated in Chapter 7.

2.5.1 Sorting Based Coarse-Fine Sensing

In Yue 2009 [48], a one-order CFD [65] is used as the secondary detector. A one-order CFD also relies on the inherent periodicity of man-made signals. In this case, the detector uses the periodicity of the mean of the signal to improve detector performance. The coarse detector uses an energy detector to sort the channels with respect to the energy in the band. Based on this estimate, the channel that is most likely to be free is scanned by the one-order CFD fine detector. The system performance is not shown for multiple channels or

using the coarse detector, only results for a single detection attempt using the one-order CFD are shown in [65]. If the analysis of Chapter 5 was extended to this architecture, using the one-order CFD as the fine detector, then the number of samples used by the coarse detector could be optimised accordingly.

In Hur 2006 [47], a wavelet based scheme was implemented. The system used wavelet transforms for the coarse section to identify spectrum opportunities. The output of the coarse section gives a tuneable resolution for the spectrum estimation without any hardware changes, a very useful feature for CR. It is not clear from this work whether the authors intended the scheme to merely sort the channels or to make decisions on occupancy, though the wideband nature of the wavelet transform estimate of the spectrum would be more suited towards the former. The fine detector used is a Temporal Signature Detection technique. Since man-made signals are generally periodic, correlating a signal with a delayed version of itself results in a peak where the period of the signal is equal to the delay used. Since it does not exhibit this feature, a signal can be detected from noise with fewer samples than an energy detector, though at an increased computational cost, especially if the period of the signal is not known in advance and must be searched for.

In Park 2006 [18], the hardware issues associated with the wavelet implementation of a coarse detector scheme were investigated. Again, the implementation uses the wavelet transform to search the spectrum in a coarse manner, with the capability of being tuned to finer resolutions. However, the main focus of [18] is on minimising the effects of the hardware components' non-linearities on the sensing results.

In Luo 2009 [50], a variant of the CF system is used. The detector uses a coarse-detector over a group of bands to find a group that might contain an unoccupied band. This method uses a coarseness in the frequency resolution, rather than in the accuracy of detection. This method allows a number of bands to be scanned with fewer operations than in traditional

CF detection schemes. The system is examined for different widths of the coarse detection scheme, but only for SNR values greater than 3dB. In addition, the SNR is a single constant value across all channels. This scheme could, however, be used to increase the speed of the detectors used here, as it identifies a group of bands that appear to contain a spectrum opportunity. If the HCFD from Chapter 7 was then used on this group of bands even greater performance gains may be realised. One issue that might reduce the effectiveness of this scheme is the assumption of equal power in each band. If there is a signal with a very large SNR in one of the groups of bands then it will increase the average SNR of the bands, such that the coarse detector decides that the band is not likely to contain any opportunies, even if some of the bands are free.

2.5.2 Deciding Based Coarse-Fine Sensing

In Maleki 2010 [51], a CF system using a cyclostationary feature detector as the fine detector is examined. An energy detector scans the channel first and, if the energy detector does not declare the channel occupied, the cyclostationary feature detector scans the channel. If either detector declares the channel occupied the next channel is then scanned. Unlike other work, the authors treat the system as a single detector and attempt to find the overall probability of missed detection and false alarm. The problem is formulated as an attempt to maximise the probability of detection, given a minimum required probability of false alarm. By deriving expressions for the thresholds of both the coarse and fine detectors enabling the combined detector to have a minimum P_{md} for a specific P_{fa}, the authors show that the optimisation can be easily performed. Simulation results show that the detector has, on average, both a reduced mean detection time and a reduced P_{md}, when compared to a cyclostationary detector alone. The system was not, however, anal-

ysed for environments where the SNR is not constant across the bands, as is studied in Chapter 6. It is envisaged that further optimisation would be possible and that the system performance would be even better, when compared to other detector types.

In Ejaz 2012 [82], the detector architecture in [51] is added to a matched filter to create an "iDetection" scheme. If the CR is aware of the parameters of the PU in the band then a matched filter can be used as it is optimum. Knowing the parameters of the PU is difficult, however, and without this knowledge no net gains accrue. In addition, the CR would not be able to detect the presence of other CRs in the channel using the matched filter, if they are of a different signal type to the primary user.

In Zamat 2008 [52], a dedicated receiver is investigated with a CF architecture. The main CR receiver does not sense the channel in this scheme, rather a dedicated receiver is used to improve performance. The system uses a constantly updating LUT to decide the channels to be scanned. This LUT is updated via the coarse sensing stage and is used to decide which of the channels are to be scanned and which are probably occupied. The coarse detector architecture is varied, using a matched filter if the signal characteristics are known. If the signal characteristics are not known, the system uses an energy detector based FFT method to generate an estimate of the spectrum. This FFT method allows the spectrum estimate to be generated in parallel for all bands. The FFT method is also quick, efficient and tuneable. The results of [52] indicate that the dedicated sensing receiver system allows significantly reduced sensing times. The dedicated sensing receiver system could begin transmitting after 50.1 ms, compared with 5.5 s for traditional architectures. These results were, however, based on a single theoretical situation and the actual improvement of the system, on average, is not derived.

In Zhang 2010 [84], using CF sensing to reduce the effects of noise uncertainty is investigated. By using a double threshold energy detector as the coarse detector it is possible to

reduce the impact of noise uncertainty. The proposed detector is more robust to noise uncertainty than an energy detector as well as being quicker than a cyclostationary detector. In addition, it has a higher probability of detection than the energy detector, especially at low SNR.

In Geethu 2012 [83], a CF system with a Covariance Absolute Value (CAV) detector for the fine detector is investigated. The CAV fine detector is based on the detector of [75]. This detector differs from conventional CF detector schemes by enabling the coarse detector to declare a band free. This ability requires knowledge of the SNR, which the detector does not estimate. If an estimate of the SNR is available, then the performance increase is significant. In addition, this architecture requires less sensing time than conventional schemes (e.g. CFDs).

In all of the CF systems discussed here the SNR is kept constant across the bands. To the author's best knowledge, there have been no CF studies published that investigate the performance of a system where the SNR is taken from a possible set of values. This situation is likely, in practice, to be more realistic and forms the basis for the work in this thesis.

In addition, few works [49–51] attempt to optimise the coarse section of the detector. Usually, the coarse detector is simply specified as an estimator using an energy detector, or some similar scheme, and the parameters of the coarse detector are not investigated. The question of how coarse the coarse detector should be is not answered. If the detectors were analysed using the methods in Chapter 5 and Chapter 6 for an SNR distribution then it is envisaged that they could be optimised and further performance gains realised. Unfortunately, it is difficult to compare the schemes as each work makes different assumptions about the prevailing channel conditions.

2.6 Cognitive Radio Test-beds

A number of test beds have been set up by research institutions to test the theoretical designs in real environments. In the Centre for Telecommunications Value-Chain Research (CTVR) in Trinity College, Dublin, the IRIS system [85, 86] is available for EU researchers. The IRIS system uses a Universal Software Radio Peripheral (USRP) as the radio front-end [87]. The USRP is the most common radio front-end in use and all the systems considered here use it. The IRIS system is used for all practical work in this thesis, primarily in Chapter 4 and Chapter 7.

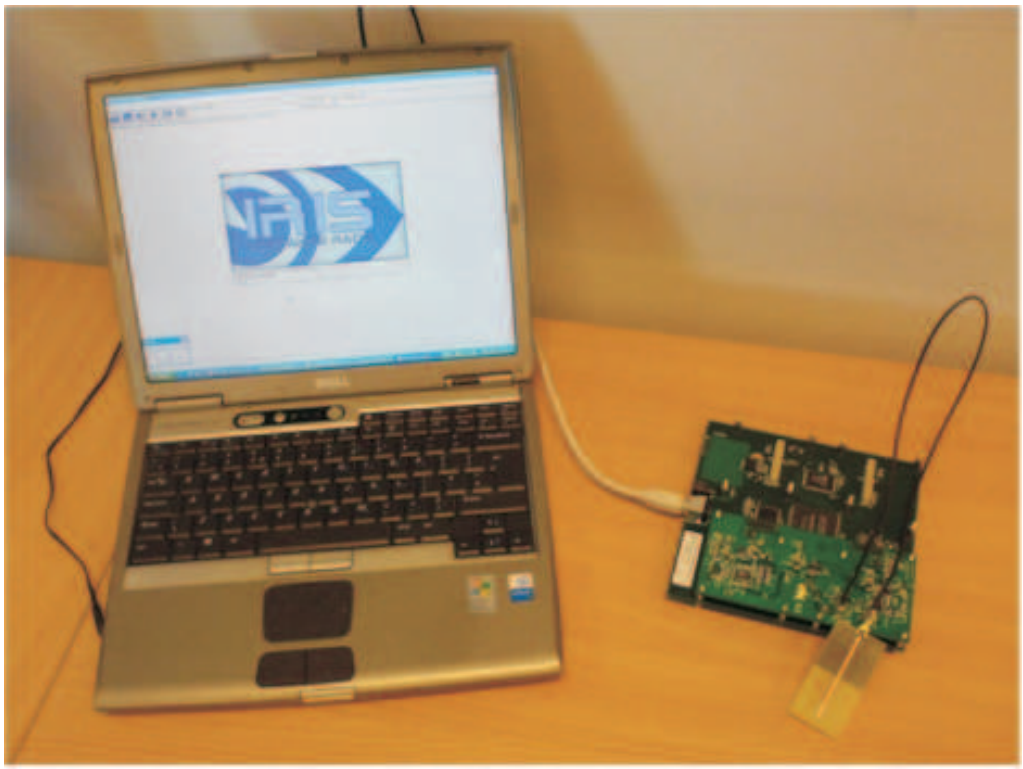

Figure 2.9: IRIS System with USRP front-end

Other test-beds exist that use similar systems such as GNU radio [88] and the Berkeley

Emulation Engine 2 (BEE2) [89–92]. GNU radio differs from the IRIS system in terms of the scope of previous implementations. GNU radio has had a long history (it commenced in 2001), many test-beds use it and many systems have been implemented on it.

In [91, 92], the BEE2 system is used to test practical issues for CR applications. The need for a practical implementation in CR systems is argued. The effects of frequency offset and noise uncertainty are investigated. The modification to the SNR walls due to FFT length is shown, where a longer FFT can have a lower SNR wall than a shorter FFT. In addition, practical results for the gain possible from collaborative sensing show that, for the signal environment used for the test, five collaborating radios improved the detection rate to 97%, compared to 63% for a single radio. The impact of spatial separation on collaborative sensing is also shown, since if the radios had experienced independent fading, the detection rate would be approximately 99%.

2.7 SNR Estimation

One of the contributions of this work is the ability to estimate a SNR distribution accurately. SNR estimation has been studied in several papers [93–96]. In Clarke 1980 [96], a closed form solution for the distribution of the SNR, when estimated using an energy detector, is derived. It shows that, for a small number of samples, it is difficult to get an accurate estimate for the SNR, especially if the SNR is low. It is recommended that the energy detector is not used for SNR estimation. In Chapter 4 an energy detector is used for estimating the SNR; however, there are a few key differences in how the energy detector is used. Firstly, a single value of SNR is not expected, rather an SNR distribution will be generated that will describe the possible range of values for the SNR. Consequently a small variance in the result is generally tolerable. Secondly, from simulation

results in Chapter 5 and Chapter 6, the systems that are being optimised are relatively robust with little performance loss for slight errors in the SNR distribution estimate. Finally, there will be a large number of samples for each estimate generated, which will reduce the variance of the output.

In Matricciani 2011 [93], an SNR distribution is estimated from Bit-Error Rate (BER) measurements in Binary Phase-Shift Keying (BPSK), Quadrature Phase-Shift Keying (QPSK) and Quadrature Amplitude Modulation (QAM) systems. The SNR at the receiver is shown to be lognormal in shape, with a mean and variance proportional to the error rate. This method works well for a sufficient number of errors, reporting SNR errors of less than 1dB for a wide range of SNR values. However, only positive SNR values are analysed. This is an inherent feature of this system, the BER measurements require data transmission. At very low SNR values the BER would be sufficiently high to render transmission too lossy to be reliable. Although for the applications envisioned in [93] this is not an issue, for CR DSA it is. Indeed, there is no way for the CR to estimate BER without knowing the system parameters. In network co-ordination, where a CN is attempting to control parameter(s), such as transmission power, to meet a certain set of requirements, this method of estimation could prove useful.

If knowing the system parameters is impractical, then non-data-aided SNR estimation must be used. In Wiesel 2002 [94], a system using BPSK is investigated, though M-ary Phase-Shift Keying (M-PSK) is stated to be easily derived from the work. The estimator architecture requires estimation of the underlying symbol PDF, rather than explicit symbol decisions. The detector allows performance close to the Cramer-Rao bound, the lower bound on the error under certain conditions. The SNR is not varied below 0dB, therefore it is difficult to ascertain the performance at the very low SNR values $\approx$ -21dB that CRs are required to detect.

In Pauluzzi 2000 [95], multiple SNR estimation techniques are compared. The metric used to judge the accuracy is the Mean Square Error (MSE). In Chapter 4 an equivalent metric, the Mean Integrated Square Error ($MISE$) is used, the difference being that the $MISE$ is used for distributions and the MSE for single values. As a single value is being estimated in [95], the results cannot be easily compared to those derived in this work. However, the work can be compared to the idealised detector, which has no cost and is completely accurate. In Chapter 4 it will be shown that the SNR estimator is both accurate and requires very little extra computation.

As with the CF detector, there is currently a lack of comparable work. The papers examined [93–96] do not extend their analysis below 0dB. In CR applications the SNR will often be significantly lower.

2.8 Conclusion

CRs have great potential to reduce the current issue of spectrum shortage but the technical challenges of DSA still need to be solved. CF sensing allows some significant performance gains over more traditional architectures. Current CF architectures are not optimised for a distribution of SNR values and this does not allow the detectors to operate at their peak efficiency. Methods to account for this SNR distribution are the main contribution of this work and most of the architectures studied in Section 2.5 would likely benefit from this analysis. Some of the topics covered here, for example energy detectors and fading, will now be covered in greater detail in Chapter 3.

3

Basic Theory

3.1 Introduction

In this chapter some relevant background signal detection theory is reviewed. The binary hypothesis test and the energy detector are fundamentals of signal detection and are used extensively throughout the literature. Here, they are introduced and discussed within the established bounds of this thesis. Effects such as noise uncertainty and fading are also considered and their effects on the signal environment are shown.

A method for modelling the transitions within a multi-state system, namely Markov Chain analysis, is also reviewed. Markov Chains are used in this work to model the effects of channel sorting in Chapter 5, and to derive related expressions for some important system properties. Finally, the IRIS test-bed is described. The fundamentals of the IRIS system and how radios can be designed using it are discussed. The IRIS system is used throughout this thesis for implementation results.

3.2 Binary Hypothesis Test

When attempting to determine the occupancy of a channel, there are two possible hypotheses. The channel can be unoccupied, denoted by hypothesis H_0, or occupied, denoted H_1. The detector estimates the channel occupancy and declares it occupied, $D(1)$, or unoccupied, $D(0)$, by use of a binary hypothesis test. There are four possible outcomes for this test.

P_d is the probability of detection, an event that occurs when the detector correctly declares the channel occupied and is given by:

$$P_d = P(D(1)|H_1), \tag{3.1}$$

where $P(x)$ denotes the probability of event x occurring.

P_{md} is the probability of missed detection, where missed detection occurs when the detector incorrectly declares the channel free and is given by:

$$P_{md} = P(D(0)|H_1). \tag{3.2}$$

P_{fs} is the probability of detecting a free space, an event that occurs when the detector

correctly declares the channel unoccupied and is given by:

$$P_{fs} = P(D(0)|H_0). \tag{3.3}$$

P_{fa} is the probability of false alarm, where a false alarm occurs when the detector incorrectly declares the channel occupied and is given by:

$$P_{fa} = P(D(1)|H_0). \tag{3.4}$$

In addition,

$$P_d + P_{md} = 1 \tag{3.5}$$

and

$$P_{fs} + P_{fa} = 1. \tag{3.6}$$

This binary hypothesis test is used for determining the channel occupancy in all of the detectors investigated here.

3.3 Energy Detectors

An energy detector operates by generating an estimate of the energy in the band. This estimate can then be used to decide if a signal is present, or not, in the band. Here we consider only a single threshold energy detector, such that the Binary Hypothesis test is applicable.

An energy detector generates an estimate of the energy in the band by finding the aver-

age energy in N instantaneous samples, $x(n)$, $1 \leq n \leq N$, assuming a 1Ω reference resistor and the noise is AWGN. Note that the following equations assume that the noise samples are independent, for example interference that is changing significantly slower than the sample rate, if this is not the case then the AWGN assumption is not valid. The characteristic equation of the detector is given by:

$$Y = \sum_{n=1}^{N} \mid x(n) \mid^2 . \tag{3.7}$$

The energy content of the band will, in practice, be non-zero even in the absence of a target man-made signal. In the absence of the target signal, this energy is due to the noise power in the channel. If a signal is present in the band, then the energy in the band will be increased by the energy of the signal, in addition to the noise power already present.

The n^{th} sample of the received signal, $x(n)$, is given by

$$x(n) = \begin{cases} v(n) & H_0 \\ v(n) + s(n) & H_1 \end{cases} , \tag{3.8}$$

where $v(n)$ is the n^{th} noise sample and $s(n)$ is the n^{th} signal sample.

The input to the energy detector is scaled by an estimate of the noise variance, σ_n^2 [54]. The noise is assumed to be zero mean, with a Gaussian distribution, therefore the noise samples are now zero mean, unit variance or :

$$v(n) \sim \mathcal{N}(0, 1). \tag{3.9}$$

The signal can also be considered a zero mean Gaussian variable with a variance of σ_s^2. After scaling, the samples are zero mean with a variance of σ_s^2/σ_n^2, denoted here as γ, or:

$$s(n) \sim \mathcal{N}(0, \gamma), \tag{3.10}$$

where γ is the SNR of the signal (in the linear scale).

Under H_0, the output of the energy detector is the sum of the squares of N zero mean, unit variance, Gaussian distributed variables. For large N this can be approximated by a Gaussian distribution with a mean of N and variance of $\sqrt{2N}$ [97] .

The output of the energy detector under H_1 for large N can be approximated by a Gaussian distribution with a mean of $N(1+\gamma)$ and variance of $\sqrt{2N(1+2\gamma)}$ [54]:

$$Y \sim \begin{cases} \mathcal{N}(N, \sqrt{2N}) & H_0 \\ \mathcal{N}(N(1+\gamma), \sqrt{2N(1+2\gamma)}) & H_1 \end{cases}. \tag{3.11}$$

To decide between H_0 and H_1 requires setting a threshold, λ. The binary hypothesis test probabilities, (3.1)-(3.4), become

$$P_d = P(Y \geq \lambda | H_1), \tag{3.12}$$

$$P_{md} = P(Y < \lambda | H_1), \tag{3.13}$$

$$P_{fs} = P(Y < \lambda | H_0), \tag{3.14}$$

$$P_{fa} = P(Y \geq \lambda | H_0). \tag{3.15}$$

In Fig. 3.1, the PDFs of H_0 and H_1, along with the relevant λ, are plotted for $N = 10,000$,

$\gamma = -15dB$ and $P_{fa} = 0.1$.

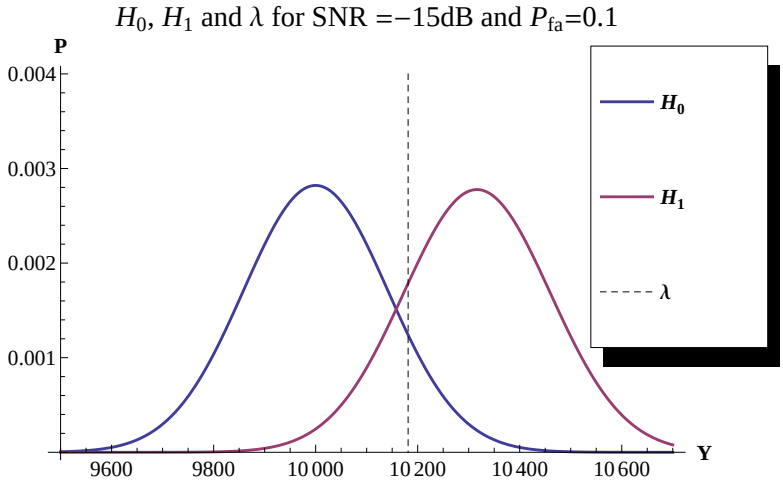

Figure 3.1: PDF's of Y under H_0 and H_1 for SNR=-15dB and P_{fa}=0.1

P_{fs} is a function of N and λ and is equal to the Cumulative Distribution Function (CDF) of a Gaussian variable with a mean of N and variance of $\sqrt{2N}$ evaluated at λ. Thus:

$$P_{fs} = CDF[\mathcal{N}(N, \sqrt{2N}), \lambda]. \tag{3.16}$$

P_d is a function of N, λ and γ and is equal to the Complementary Cumulative Distribution Function (CCDF) of a Gaussian variable with a mean of $N(1+\gamma)$ and variance of $\sqrt{2N(1+2\gamma)}$ evaluated at λ. The P_d of the system associated with Fig. 3.1 is approximately 82.6%.

$$P_d = CCDF[\mathcal{N}(N(1+SNR), \sqrt{N(1+SNR)}), \lambda]. \tag{3.17}$$

It is possible, then, to solve for the values of N and λ that give the required P_d and P_{fs} for a SNR of γ. In practice, it is more usual to give them as a function of P_{fa} and P_{md} .

Therefore:

$$N = \frac{2[Q^{-1}(P_{fa}) - Q^{-1}(1 - P_{md})(1 + \gamma)]^2}{\gamma^2} \tag{3.18}$$

where $Q^{-1}(\cdot)$ is the Inverse Cumulative Distribution Function (ICDF) of the normal distribution and

$$\lambda = ICDF[\mathcal{N}(N, \sqrt{N}), P_{fa}]. \tag{3.19}$$

For a P_{fa} and P_{md} of 0.1, the required values for N and λ for varying SNR are shown in Fig. 3.2.

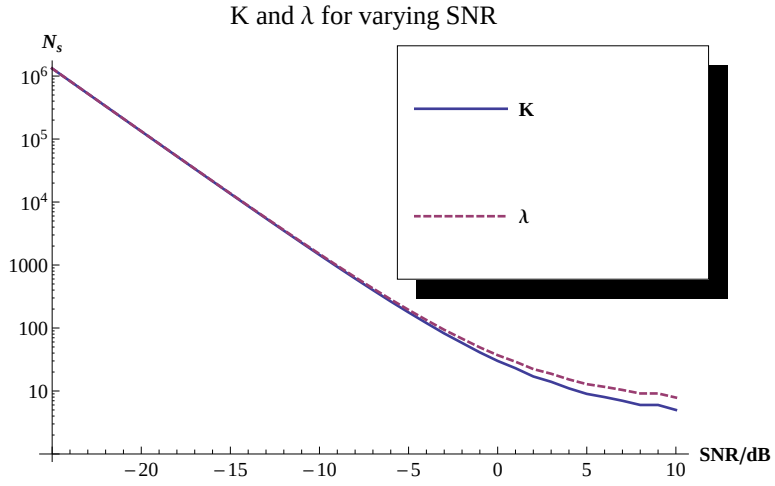

Figure 3.2: Required N and λ for P_{fa}=0.1 and P_{md}=0.1 Under Varying SNR

Note the large rise in the required number of samples as the SNR decreases, indeed $N \propto \frac{1}{\gamma^2}$. This leads to inefficiencies in sensing that will be discussed in later chapters.

3.4 Energy Detector Outputs under Noise Uncertainty

Thus far in the analysis of the energy detector it has been assumed that the system has perfect knowledge of the noise variance, σ_n^2. This variance is used to scale the input of the energy detector to give a unit variance input when no signal is present. However, if the estimate of the variance, $\hat{\sigma}_n^2$, is incorrect, then the energy detector input will be scaled incorrectly.

Consequently, ρ is defined as the noise-uncertainty, or the relative inaccuracy of the estimate, and is given by [42]:

$$\rho = \frac{\hat{\sigma}_n^2}{\sigma_n^2}. \tag{3.20}$$

For an unoccupied channel, the distribution of the input to the energy detector, after normalisation, has variance ρ:

$$v(n) \sim \mathcal{N}(0, \rho). \tag{3.21}$$

The energy detector output:

$$Y_n = \sum_{n=1}^{N} v(n)^2, \tag{3.22}$$

is distributed according to,

$$Y_{n,\rho} \sim \sum_{n=1}^{N} \mathcal{N}(0, \rho)^2, \tag{3.23}$$

For an occupied channel, the n^{th} input to the energy detector, $x(n)$, after normalisation, is the sum of two zero mean Gaussian variables, namely a noise variable with variance ρ

and a signal of variance $\frac{\gamma}{\rho}$.

Therefore, $x(n)$ is the sum of two Gaussian variables and has distribution:

$$x(n) \sim \mathcal{N}(0, \rho + \frac{\gamma}{\rho}). \tag{3.24}$$

The energy detector output is then distributed according to:

$$Y_{n,\rho} \sim \sum_{n=1}^{N} \mathcal{N}(0, \rho + \frac{\gamma}{\rho})^2. \tag{3.25}$$

Clearly, the detector will find it more difficult to detect the signals in this case. In (2.3) the total number of samples required for reliable detection is shown. This equation shows that the number of samples approaches infinity as the noise uncertainty reaches a certain value, which depends on the SNR of the signal. Noise uncertainty is a severe problem in energy detector based architectures. If noise uncertainty is present, then reliable detection becomes significantly more difficult.

3.5 Fading Channels

If the system is subjected to fading then, instead of a single static SNR value, the SNR now has a probability distribution associated with its instantaneous value. This could be caused by multipath effects or shadowing [41]. There are many fading types with corresponding distributions, such as Rayleigh, Nakagami-m and Ricean Fading. In addition to the fading distribution type, there are also two different classifications of fading types, namely fast and slow fading.

Fast fading occurs when the channel impulse response changes sufficiently quickly such that the SNR is varying during each detection attempt and not just between detection

attempts. For an energy detector, fast fading does not cause significant problems. The inherent averaging performed by the energy detector means that the performance will be close to that for a system with an SNR equal to the average SNR. Thus, fast fading is not considered here.

Slow fading occurs when the channel impulse response changes slowly, such that the SNR is constant within each detection attempt but may vary between detection attempts. This form of fading is considered in this work.

For example, with Rayleigh fading [55], the PDF of the SNR, γ, with an average SNR of $\bar{\gamma}$, $f_{fad}(\gamma, \bar{\gamma})$ is:

$$f_{fad}(\gamma, \bar{\gamma}) = \frac{1}{\bar{\gamma}} e^{-\gamma/\bar{\gamma}}; \gamma \geq 0. \tag{3.26}$$

Distributions for Ricean and other channels, as well as methods for finding the probabilities of missed detection for a signal under fading, can be found in [55].

For an SNR distribution, like that which will be considered in this work, fading changes the distribution.

Replacing the single value of $\bar{\gamma}$ in (3.26) with the user defined SNR distribution, $f_{SNR}(\bar{\gamma})$, and averaging, the new SNR distribution under fading can be calculated by

$$f_{SNRfaded}(\gamma) = \sum_{\bar{\gamma}} f_{SNR}(\bar{\gamma}) f_{fad}(\gamma, \bar{\gamma}). \tag{3.27}$$

Fading can significantly reduce the performance of a system. An energy detector subjected to slow Rayleigh fading with an average SNR of -21dB requires approximately 4.8M samples to detect a signal with a P_{md} of 10% and a P_{fa} of 10%, compared to 209k when fading is not present. This value can be found by averaging the P_{md} over the SNR distribution of a Rayleigh faded signal at an average SNR of -21dB [55].

3.6 Markov Chains

A Markov Chain is a model of a system of transitions between groups of states of finite size. Any system that is memoryless (i.e. the next transition depends only on the current state), and has a countable number of possible states, can be represented by a Markov Chain [98] .

A simple example is a system such as a switch with three states, “on” (1), “off” (0) and “broken”. When in the “off” state, there is a probability of turning “on” of P_{on} and a probability of staying “off” of $1 - P_{on}$. When in the “on” state, there is a probability that the switch will turned “off” of P_{off}, a probability that it will break of P_{break} and a probability that it will stay “on” of $1 - P_{off} - P_{break}$. If the system enters the “broken” state then it cannot leave, this is termed an absorbing state. The “on” and “off” states are transient states as the system, given enough time, will always leave them.

Traditionally, Markov Chains are drawn as a directed graph and the graph for this example is displayed in Fig. 3.3.

There exists a transitional probability matrix, $\mathbb{Q}$, associated with every Markov Chain, which gives the transition probabilities between all the transient states. The $\mathbb{Q}$ matrix for this example system is:

$$\mathbb{Q} = \begin{pmatrix} & off & on \\ off & 1 - P_{on} & P_{on} \\ on & P_{off} & 1 - P_{off} - P_{break} \end{pmatrix}. \tag{3.28}$$

For the absorbing states there is a matrix, $\mathbb{R}$, which gives the transition probabilities between the transient states and the absorbing states. The $\mathbb{R}$ matrix for this system is:

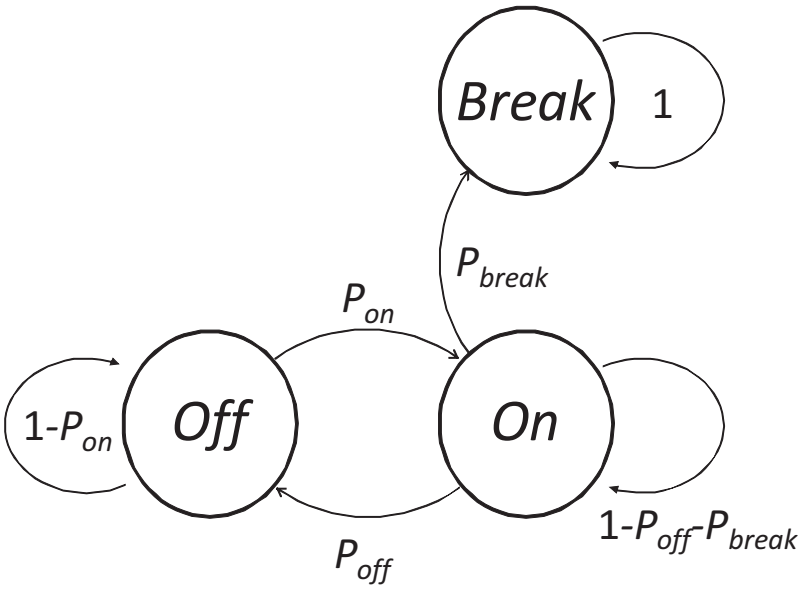

Figure 3.3: Markov Chain Represented in Directed Graph Format

$$\mathbb{R} = \begin{pmatrix} 0 \\ P_{break} \end{pmatrix}. \tag{3.29}$$

The fundamental matrix, $\mathbb{N}$, is given by:

$$\mathbb{N} = (\mathbb{I} - \mathbb{Q})^{-1}, \tag{3.30}$$

where $\mathbb{I}$ is an identity matrix of the same order as $\mathbb{Q}$ and $\mathbb{X}^{-1}$ denotes the inverse of the matrix $\mathbb{X}$.

Using these matrices it is possible to derive expressions for some of the properties of the system. This will be used in Chapter 5 to analyse the CSFD architecture. In particular, the markov approach is used to model the effects of sorted channels in Chapter 5, and to derive expressions for some of the properties, such as the average number of detection

attempts and the probabilities of interference, of those systems.

3.7 Probability of Channel Changing State During Detection Attempt

In some of the work studied in Chapter 2 [43, 45], the degradation in performance caused by time varying occupancies was considered. There is a non-zero probability that, for a time varying channel, the occupancy will change during a detection attempt and this will cause a reduction in the accuracy of the detector. This was shown in Fig. 2.6. The magnitude of this change depends on the length of the detection window and also the rate at which the channel changes occupancy. An exponential distribution is commonly used to model the 'ON'-"OFF" time of transmissions [56, 57] and this assumption will also be used here.

For a channel which is "ON", the PDF at a time t_{sense}, $f_{ON}(t_{sense})$, is given by:

$$f_{ON}(t_{sense}) = \frac{e^{-t_{sense}/t_{ON}}}{2t_{ON}}, \tag{3.31}$$

where t_{ON} is the average length of a transmission for the "ON" state and t_{sense} is the length of time required by the detector to make a decision.

Similarly, the "OFF" PDF at a time t_{sense}, $f_{off}(t_{sense})$, is given by:

$$f_{off}(t_{sense}) = \frac{e^{-t_{sense}/t_{OFF}}}{2t_{OFF}}, \tag{3.32}$$

where t_{OFF} is the average time between transmissions or the "OFF" state.

Thus, it is possible to determine the probability that the channel will change its occupancy

state within a certain time period, t_{sense}. However, this probability can only be calculated when the start (or end) time of transmission is known. If this is not known, then the current value of t_{sense} is not available.

Instead, consider the situation after t seconds, the PDF at a time t is:

$$f_{ON}(t) = \frac{e^{-t/t_{ON}}}{2t_{ON}}. \tag{3.33}$$

and, consequently, t_{sense} seconds later f_{ON} is:

$$f_{ON}(t + t_{sense}) = \frac{e^{-(t+t_{sense}))/t_{ON}}}{2t_{ON}}. \tag{3.34}$$

However, if the channel is still occupied after t seconds this influences the probability of staying 'ON" after $t + t_{sense}$ seconds. The probability of being occupied after $t + t_{sense}$, given that after t seconds the channel remains occupied, is:

$$P_{ON}(t + t_{sense} \mid t) = \frac{P_{ON}(t + t_{sense})}{P_{ON}(t)}, \tag{3.35}$$

where $P_{ON}(t + t_{sense})$ and $P_{ON}(t)$ are given by the CCDF of the exponential distribution at $t + t_{sense}$ and t respectively, or:

$$P_{ON}(t + t_{sense}) = e^{-\frac{t+t_{sense}}{t_{ON}}}, \tag{3.36}$$

and

$$P_{ON}(t) = e^{-\frac{t}{t_{ON}}}. \tag{3.37}$$

Equation (3.35) becomes:

$$P_{ON}(t + t_{sense} \mid t) = e^{-\frac{t_{sense}}{t_{ON}}}. \tag{3.38}$$

Similarly, for the "OFF" state:

$$P_{OFF}(t + t_{sense} \mid t) = e^{-\frac{t_{sense}}{t_{OFF}}}. \tag{3.39}$$

The probability of a change occurring in the occupancy within t_{sense} seconds is, therefore:

$$P_{change} = \begin{cases} 1 - e^{-\frac{t_{sense}}{t_{OFF}}} & \text{“}OFF\text{”} \\ 1 - e^{-\frac{t_{sense}}{t_{ON}}} & \text{“}ON\text{”} \end{cases}. \tag{3.40}$$

If the result of (3.40) is sufficiently high (e.g. $P_{change} \geq 0.05$, for example), then the effects of a changing occupancy will need to be accounted for. This analysis will be performed on the CSFD and CDFD architectures in sections 5.5 and 6.3.2, respectively.

3.8 IRIS Architecture

The physical implementations for this thesis were implemented on the IRIS system. The IRIS system is implemented on a CR/SDR test-bed hosted in Trinity College, Dublin. The test-bed is under the control of the CREW project and is available to European Union researchers upon request.

IRIS is a software radio architecture which has been developed by CTVR in Trinity. It is written in C++ and implements a fully reconfigurable radio system. There is a minimal hardware front end, in the form of a USRP, that provides a receiver and transmitter for the radio. Each front-end is paired with a PC which runs the rest of the radio implementation in software. The IRIS architecture allows for reconfigurable architectures to be run in real

time on multiple nodes. There are a number of nodes available in Trinity, nine of these nodes were available for the work in this thesis, and these can be used to create sample environments or networks.

In Fig. 3.4 the IRIS architecture is shown[i].

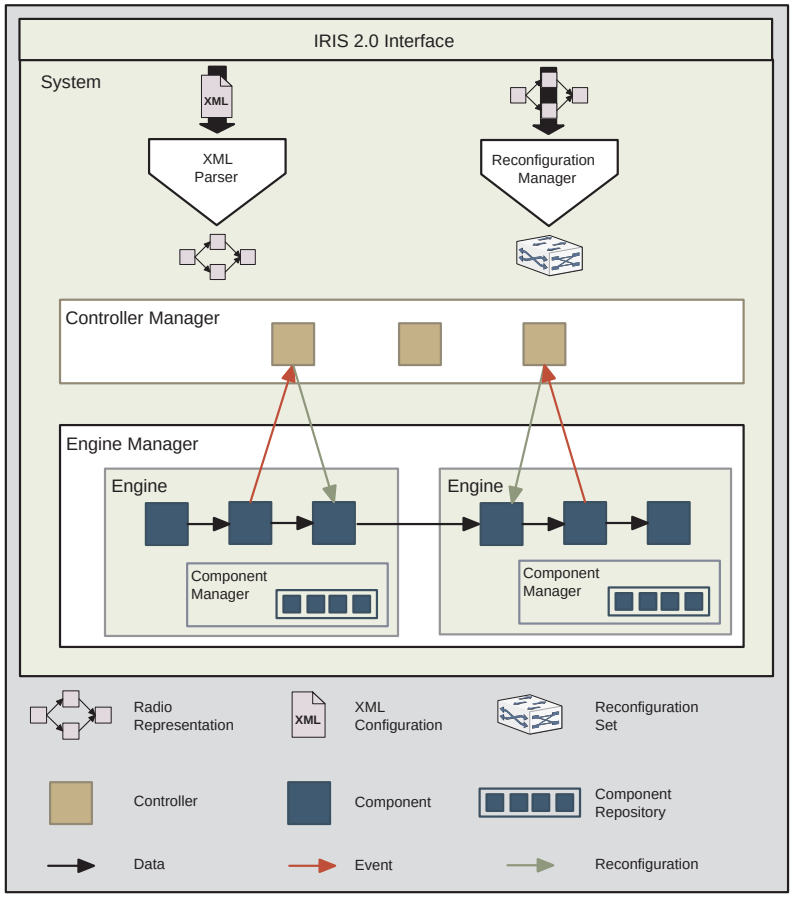

Figure 3.4: IRIS Architecture

At the core of the all IRIS radio implementations is an .xml file that sets the initial configuration of the system. This includes the components to be used, the controllers that allow automatic reconfiguration and the links between them. It is possible to reconfigure the

[i]reproduced with permission

radio by editing this .xml file but this is not ideal as it requires direct input from the user. Each of the components has a single task that the radio can be reconfigured to use, such as low pass filtering or Orthogonal Frequency-Division Multiplexing (OFDM) modulation. Upon initialisation, the radio loads the components from a library controlled by the component manager. Each of the components can be also reconfigured. For example, the number of frequency bins in the OFDM modulation component can be changed. The components can also be reconfigured manually by editing the .xml file.

Components are available as stack components or data flow process network (PN) components. The stack components allow a bi-directional data flow and are designed primarily with higher layers of the OSI seven layer model in mind. The components used in this work are designed exclusively as PN components with strict uni-directional data flow between components.

Around the components, controllers are set up that can automatically reconfigure the radio. These controllers require triggers sent by the components and allow the radio to adapt without user input. The controllers can then pass the new parameters to the components.

All of the components and controllers are implemented in software on the radio. Indeed, if no transmission or reception of signals is required by the system, such as in debugging, then the USRPs are not required and the radio can be run entirely in software.

Although IRIS comes with a number of components already designed, some specialised components and controllers have been generated for this work. Since the time of writing the IRIS code base has been updated, making these components obsolete. The original plan of including these components in the IRIS libraries cannot now occur.

A good example of the IRIS system in action is from a demonstration in DYSPAN 2007 [99], where the IRIS system was performing system rendezvous. The transmitting radio changed the frequency of transmission and the receiver lost its lock. The controller in the

second system noticed this event and reconfigured the front end to a new frequency and attempted to find the signal again. In this case a cyclostationary feature detector attempted to find the cyclic prefix and, thus, the new operating frequency. Once the signal had been found, the controller again reconfigured the radio. It changed the operating frequency to the correct value and restarted whatever processing was being done on the signal received. In the case of the demonstration radios there was a song being sent by the transmitting radio's PC and played by the receiving one. The video of the demonstration can be found at [99] and is based on [61].

This reconfigurability will be used to allow the architectures in this work to be tested. Where learning or adaptation is required, the controllers can implement the learning and reconfigure the radios as required.

3.9 Conclusion

The relevant background theory underpinning signal detection has been introduced. In addition, the basics of Markov Chains and issues with non-stationary channel occupancies have also been discussed. The IRIS system has been outlined and, with this background information, the main work of the thesis can be presented.

4

SNR PDF Estimation

4.1 Introduction

In this chapter the SNR of a signal at a receiver is investigated and the argument for using a distribution, rather than a single value, to describe the received SNR is made. By examining a typical environment, it will be made clear that the concept of an SNR distribution is valid, and that these distributions are of practical use to the designer of CR receivers. The SNR distributions derived here will allow further optimisation later in the

thesis.

To generate the estimates of the SNR distributions based on observations of the SNR, Kernel Density Estimation (KDE) is used. The basic theory of KDE is explained and two metrics for calculating the accuracy of the estimation are shown.

Three methods for generating the observations of the SNR values for use with KDE are discussed. The three methods are:

1. Analytically in advance.
2. Experimentally in advance.
3. In-line sensing.

These methods are investigated and the last, in-line sensing, is shown to be the most promising.

In-line sensing uses an energy detector to provide estimates of the SNR. Three issues that may cause inaccuracy in the estimated SNR values are investigated, namely noise uncertainty, the threshold of the energy detector and occupancy of the channel. Finally, the accuracy of the estimate of the SNR distribution will be investigated, both in simulations and by practical implementation on the IRIS system.

4.1.1 Signal-to-Noise-Ratio

The SNR of a band limited signal, γ, is the ratio of the received power of a signal, P_s, to the power of the noise in that band, P_n or:

$$\gamma = \frac{P_s}{P_n}. \tag{4.1}$$

For example, if there is a signal present in a band with a power at the receiver of 1W and there is 10mW of noise in the band, then the SNR is 100. SNR is usually expressed in dB, where:

$$\gamma_{dB} = 10log_{10}\gamma. \tag{4.2}$$

In this case the SNR in dB, γ_{dB}, would be 20dB.

The total noise power in a band is dependent on the bandwidth of the receiver and the average noise power per Hz. The average noise power per Hz can be expressed in W/Hz, and, thus, is independent of the bandwidth of the channel. To find the total noise power in any band it is then simply a case of multiplying the average noise power per Hz by the bandwidth. If the front-end filter has a bandwidth of 8MHz, and the noise power is 10nW/Hz, then the total noise power is 80mW. In this work a channel width of 8MHz is assumed. This is the bandwidth of the Irish DTV channels and will be used to estimate the SNR for signals later in this chapter.

Typically, the average noise power is approximately constant across a group of bands, provided there is no narrowband noise, e.g. no leakage from adjacent channels. The signal power, however, will vary significantly due to distance and transmit power. If a signal is transmitted from a station with a transmit power of P_T, then the signal power at a receiver a distance d away, P_R, satisfies:

$$P_R \propto \frac{P_T}{d^n}, \tag{4.3}$$

where n, the order of d, depends on the path type. Typically, the value of n varies from 2-4, though for high path loss environments, such as indoor environments, n can be as large as 6.

In free space, where $n = 2$, the power decreases with $O(d^2)$. For some channels subjected to fading, $n = 4$ and the power decreases with $O(d^4)$. In fading environments the power at a distance d is distributed with the appropriate situation specific distribution [100]. This distribution of SNR values is due to the time-varying nature of the channel and can be caused by a large number of factors. These factors range from atmospheric effects to the motion of vehicles in the channel and can have a detrimental effect on the received signal power. The instantaneous signal power is, therefore, difficult to find and, after a short period, usually obsolete. Of greater importance are the average SNR and the SNR distribution. These parameters can be used to accurately model the channel and the performance of any system that receives signals from it [100].

Even if the signals are not subjected to fading, a group of signals in adjacent channels will have a distribution of power values. The transmitters typically have different paths, transmission powers and distances to the receiver and, thus, different signal powers at the receiver. There is, however, a distribution that describes the SNRs of all the channels. This distribution describes the signal powers of all the channels together, not the channels individually. Instantaneously the powers in the bands may not match this distribution but, as the number of observations increases, the histogram of the SNR estimates will converge towards this distribution. As in the fading distribution case, this distribution can be used to model the channel responses and the performance of a system attempting to detect a signal in the band.

In this chapter, various methods of finding this SNR distribution are discussed and the most promising method, in-line sensing, is investigated further.

4.1.2 The SNR Distribution for CR applications

When a CR attempts to find a free band, the radio is required, by current standards [3], to detect a signal at an SNR of -21dB with a P_{fa} and P_{md} of, at most, 10%. The receiver complexity required to detect such weak signals is significant, as was shown in Section 3.3. An energy detector attempting to detect signals with this SNR requires approximately 209k samples to meet these P_{fa} and P_{md} constraints. However, not all signals are at such a small SNR; some will have larger SNRs and, thus, will require less signal processing to detect, whilst still meeting the constraints on P_{fa} and P_{md}.

Consider the situation in Fig. 1.1, reproduced in 4.1 for convenience. There is a CR scanning four bands as it is attempting to detect a spectrum opportunity, or free space, in one of the four bands. The ranges at which the CR is required to detect transmissions from each of the four PUs are shown for each of the bands. Some transmitters will be closer to the CR and have a relatively higher SNR, such as in Band 1. Others will be on the edge of the detection range, such as in Band 3, and will have SNR values close to -21dB. Some transmitters are so far from the CR that it can transmit without fear of interfering, such as in Band 4.

Using current techniques, when attempting to optimise the CR, the radio typically assumes that all of the signals in the bands have the same SNR. Under the current standard, all signals have to be assumed to have a SNR of -21dB to guarantee that the CR does not interfere with any of the PUs.

This approach leads to significant inefficiencies. In contrast, CF spectrum sensing, which is investigated in further chapters, can reduce this overhead significantly, reducing the total number of samples required by 50% in some cases. Naturally, if the CR uses the minimum required processing and number of samples, then the overall system is opti-

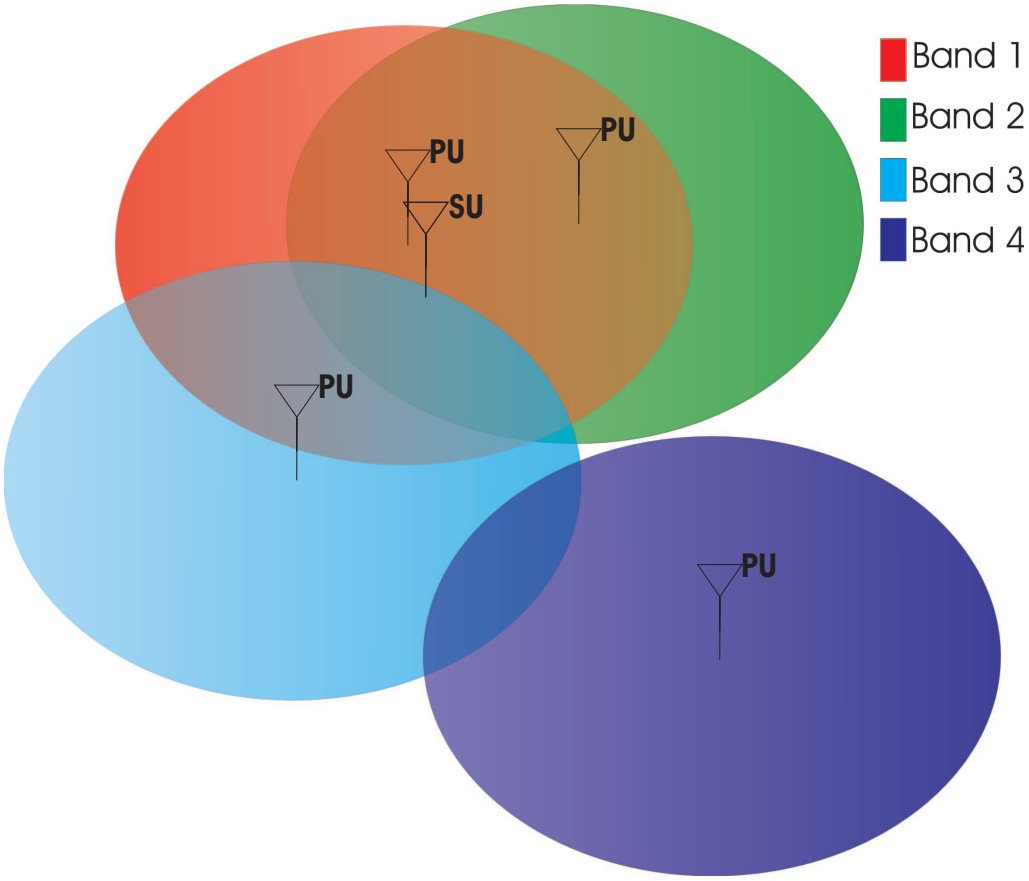

Figure 4.1: Typical Signal Environment with PUs at Different Distances to the CR

mised.

This optimisation can be performed over the SNR distribution to increase system performance, relative to a system optimised for a single SNR value. Provided that the SNR distribution is only used to increase the speed of detection, and does not change the interference probabilities, then the overall performance can be increased without any increase in interference to the PUs. This "decoupling" of the interference probabilities and SNR distribution ensures that, even if the SNR distribution is very inaccurate, only the CR's performance suffers and that the PUs remain safe from interference. Before optimising

for the SNR distribution, however, the SNR distribution must be known.

4.2 Kernel Density Estimation

In general, the SNR distribution will be estimated from a number of sample observations, N_{obs}. Although it may be possible to generate an analytical expression for the distribution, the wide range of operationing conditions that CRs could be subjected to makes this approach impractical. A histogram method could be used, and would be sufficiently accurate, provided the number of observations is sufficiently large.

A more effective method for estimating an underlying distribution from a series of observations is KDE. KDE operates in a manner similar to the histogram method but, instead of a single value, there is a kernel centred on that value. The KDE method converges on the correct distribution more quickly than the histogram method, provided the distribution is smooth and the width of the kernel is correctly specified [101].

The kernel can be any distribution having unit area. In this work a Gaussian kernel is used. The kernel has a mean of the observation value and its variance depends on the application. Fig. 4.2 shows the operation of the KDE for a Gaussian kernel with a unit variance. Five observations are shown, with the appropriate Gaussian kernel centred on each one. The estimated distribution is the sum of these distributions. The distribution needs to be divided by the number of observations to ensure the area under the resulting distribution remains equal to unity. To aid clarity, the kernels were divided by a larger number than the number of observations (1.3 times N_{obs}) in Fig. 4.2, this allows the shape of the kernels to be seen more clearly.

Selecting the correct variance for the kernel is difficult, as the distribution is not known in advance. In [102], a method for setting the variance, $\hat{h}_{opt}$, is proposed (though in [102]

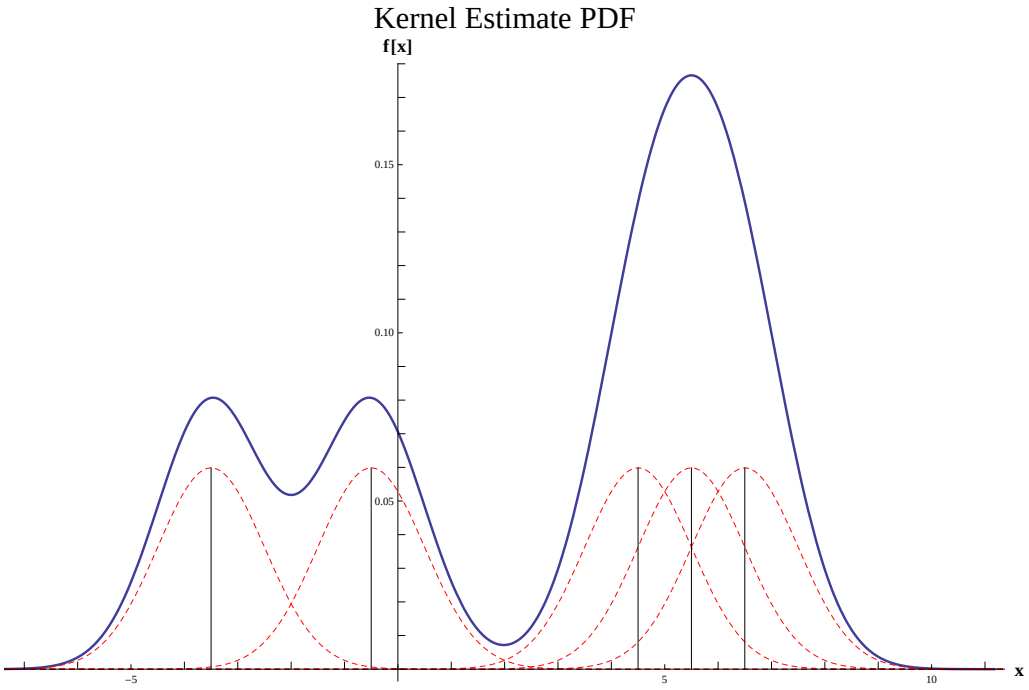

Figure 4.2: Sample PDF Based on Observations Using Kernel Density Estimation

the variance was defined as the bandwidth of the kernel) as:

$$\hat{h}_{opt} = \frac{0.9\hat{\sigma}}{N_{obs}^{\frac{1}{5}}}, \tag{4.4}$$

where $\hat{\sigma} = min(s, R/1.34)$, R is the interquartile range of the data and s is given by:

$$s = \sqrt[+]{\frac{1}{n-1}\sum_{i=1}^{N_{obs}}(x_i - \bar{x})^2}. \tag{4.5}$$

The kernel weighting function of the i^{th} estimate, $w_i(x, h)$, is :

$$w_i(x, h) = \frac{1}{\sqrt{2\pi}h} e^{-(1/2)(\frac{(x-u_i)^2}{h})}, \tag{4.6}$$

where u_i is the value of the i^{th} observation and $h = \hat{h}_{opt}$.

The PDF of the distribution at x, $f(x)$, is the mean of all the weighting functions at x or:

$$f(x) = \frac{1}{N_{obs}} \sum_{i=1}^{N_{obs}} w_i(x, h). \tag{4.7}$$

4.2.1 Accuracy of Estimated Distribution

To measure the accuracy of an estimate of a distribution, the Mean Integrated Square Error ($MISE$) can be used. The $MISE$ is the integral of the difference between the estimated distribution, $\hat{f}(x)$, and the actual distribution, $f(x)$ over the full ranges of the distributions. It is the integral of the MSE and is given by [102]:

$$MISE(\hat{f}) = \int_{-\infty}^{\infty} (\hat{f}(x) - f(x))^2 \, \mathrm{d}x. \tag{4.8}$$

This method of checking the accuracy of the system is only possible when the distribution is already known. For unknown distributions a different approach is required. Note that when the $MISE$ compares analytical distributions an expectation operator is used. In this work, sampled distributions are considered and the expectation operator omitted.

As the number of observations increases, the accuracy of the estimated distribution will, on average, increase. In addition, as the number of observations increases, the differences between the resultant distributions decrease also. For example, the difference between the distributions found after ten observations and after twenty observations will be, on average, greater than the difference between the distributions found after 100 observations and after 200 observations. This is due to the extra information and also the reduction in the effects of the outliers.

The following metric is proposed: the $MISE$ between the distribution after n observa-

tions, $f_n(x)$, and the distribution after m observations, $f_m(x)$, gives the Mean Integrated Square Change ($MISC$), or:

$$MISC(\hat{f}) = \int_{-\infty}^{\infty} (\hat{f}_n(x) - \hat{f}_m(x))^2 \, \mathrm{d}x. \tag{4.9}$$

where $m = kn$, $k \in N_0$. The $MISC$ is the measure of stability in the estimated distribution. The $MISC$ can be used to decide that the distribution has reached a sufficient level of accuracy.

Note that the multiplicative relationship between m and n is important for fair comparisons. If $m = n + k$ then, as n increased, the relative size of k would decrease and, thus, the change caused by the extra samples would not be as large. In this work $k = 2$ was chosen and the first check occurred after twenty observations. Those values were chosen to reduce the average number of observations based on the initial simulation results.

If the $MISC$ is below a user defined threshold, then the estimate of the PDF can be declared to be sufficiently accurate. The $MISC$ does not guarantee sufficient accuracy. Instead, a low $MISC$ indicates a static distribution. This, however, is a strong indicator of accuracy since a low $MISC$ would require the same outliers to be present in both sets of estimates if the estimated distribution was inaccurate.

4.3 Generating Observations of SNR Values

To use the KDE method to find a distribution, a set of SNR estimates must first be generated. Whilst estimating the SNR can be done in numerous ways, here three options are investigated. However, it will be shown that only the last, in-line measuring, is of practical use for real CR systems.

The three methods are:

1. Analytically in advance: Given transmitter locations and knowledge of the likely channel conditions, it is possible to estimate the power of the signals and, thus, the SNR. Here the distribution associated with the Irish DTV system is estimated. This method isn't practical for real implementation due to the issues discussed here, however, it is used to generate a sample distribution that is used to evaluate the detectors in later chapters.

2. Experimentally in advance: From measurements taken previously, it is possible to estimate the SNR of the signals. Here measurements are taken by a spectrum estimator of signal powers in the Irish DTV bands.

3. In-line measuring: If the system estimates the power of the signals found, it can learn the distribution over time. Here a method is derived and practical and simulation results shown.

4.3.1 Analytically in Advance: DTV Estimate

To demonstrate the difficulties in generating a distribution via analytical means, the Irish DTV system is analysed and a sample SNR distribution generated from this analysis. In Fig. 4.3 the locations of the Irish DTV transmitters are shown [103]. To find the PDF of the SNR we take the average over a range of possible locations for the CR. To maximise the number of people covered, we analyse a number of cities and large towns and, with these six locations, it is possible to cover approximately 33% of the population and a wide range of geographical locations. The transmission powers and the distances (in kilometers) from each of the chosen locations for the CR to each of the DTV transmitters are shown in Table 4.1.

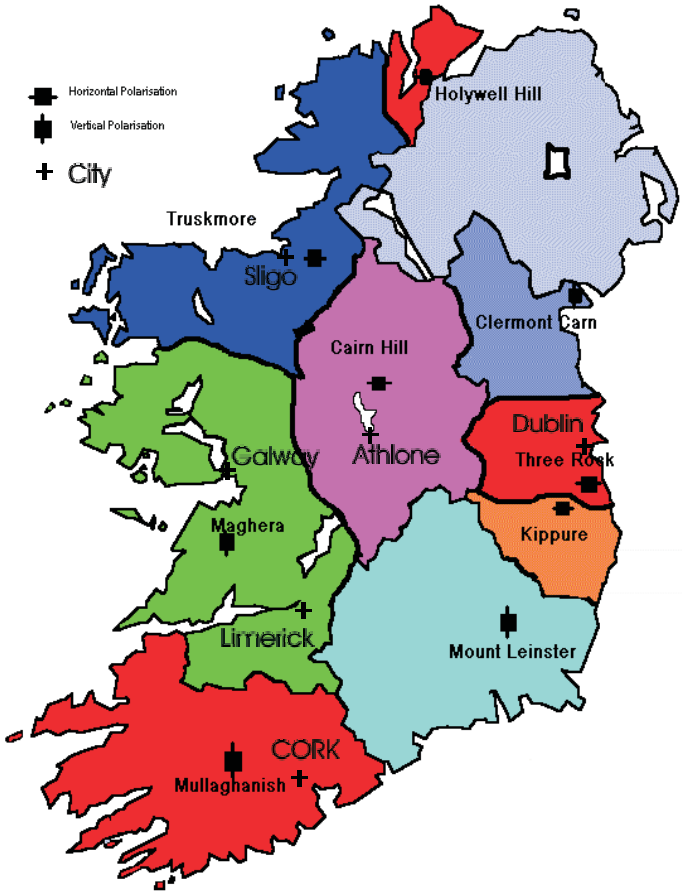

Figure 4.3: Irish DTV Locations

To find the received power at a distance d from a transmitter, a slightly modified version of Friis' equation [104] can be used:

$$P_R = P_T G_T G_R (\frac{\lambda}{4\pi d})^n, \tag{4.10}$$

where G_T and G_R are the gains of the transmitting and receiving antennae, respectively, and n is the distance-loss factor. A Okumura-Hata propogation could be used for the urban environments, however it has several parameters that would need to be estimated

Table 4.1: Locations and Distances to Urban Areas of Irish DTV Transmitters

Transmitter	P_T/W	Dublin /km	Cork /km	Limerick /km	Athlone /km	Sligo /km	Galway /km
Three Rock	50000	11	212	173	115	186	186
Kippure	50000	19	203	164	110	185	181
Greystones	10000	26	213	179	128	202	198
Clermont Carn	25000	82	282	219	129	142	200
Cairn Hill	50000	109	218	140	45	71	103
MT Leinster	50000	88	140	125	119	215	170
Dungarvan	10000	166	62	93	150	248	165
Spur Hill	10000	224	5.6	90	178	268	165
Mullaghanish	50000	245	47	82	179	257	146
Woodcock Hill	4000	178	89	5	95	176	72
Maghera	50000	198	118	41	100	161	46
Truskmore	50000	179	274	191	110	13	126
Holywell Hill	6000	198	350	272	179	106	217

and the accuracy would still not be much greater than the Friis equation.

In dB, (4.10) becomes:

$$P_{R,dB_W} = P_{T,dB_W} + G_{T,dB} + G_{R,dB} + 10nLog_{10}(\frac{\lambda}{4\pi d}). \tag{4.11}$$

Assuming a quarter wave antenna, the antenna gain is 5.14dB relative to an isotropic antenna. Thus, $G_{T,dB}$ and $G_{R,dB}$ are 5.14dB. The DTV bands are broadcast at approximately 600MHz, thus f=600MHz and $\lambda = 0.5m$. The loss exponent due to distance is given by $n = 3$; this was chosen to represent the fact that, while the signals would travel through mostly free space, they will be subjected to other effects which increase the value of n and is consistent with practical measurements taken in these bands [105].

Assuming a noise floor of -174dBm/Hz (ideal receiver at a temperature of 290 K) and 8MHz bands, there is a noise power of approximately -104dBm in each band (although if strong signals are present in adjacent bands then there will be leakage, increasing this value).

The SNR distribution is shown in Fig. 4.4. A smoothing function was applied to the dis-

tribution. The smoothing is performed by the kernel density method described in Section 4.2. In addition, the distribution used in previous work [106, 107] is also shown.

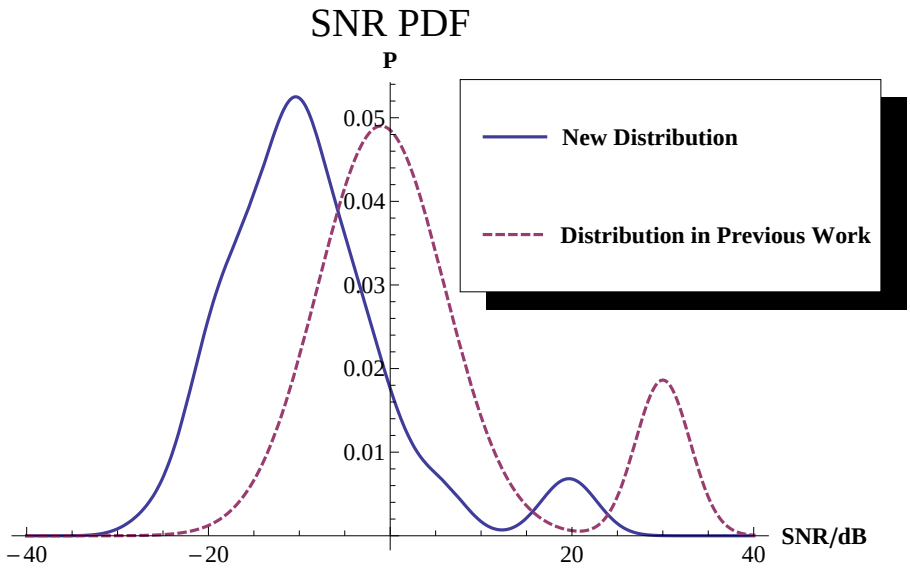

Figure 4.4: The DTV SNR PDF

Note that the distribution in [106–108] differs from that derived here. Different parameters were chosen, such as a lower frequency and a smaller value for n was used in the previous work. In addition, smoothing was performed via a moving average filtering rather the KDE method. The distribution was then matched to an appropriate Gaussian distribution. Finally, an extra city, Galway, was added to the analysis.

This difference between two distributions based on the same data shows one of the main difficulties of this method. In practice, it is very difficult to analytically determine the signal power at a distance from a transmitter with a high degree of accuracy. It was shown in [105] that the actual pathloss exponent measured can vary significantly in practice. In

addition, this method is completely inappropriate for systems where the locations and transmit powers of the PUs are not known in advance. Finally, in urban environments, where the path loss is location specific and time varying [109], an analytical solution becomes impractical.

4.3.2 Spectrum Measurements

An alternative to the analytical method presented above is to perform measurements and use them to determine the SNR PDF of a signal environment. This method solves some of the problems of the analytical method. The actual signal power can then be found directly for that location.

With the help of the CTVR in Trinity College Dublin, measurements were taken in the Irish TV bands. This set of measurements was taken outdoors in the vicinity of Mullingar, a small town in Ireland. The measurements were in the range 660MHz to 766MHz on April 07, 2012.

The equipment used for these measurements was an Anritsu MS2721B handheld spectrum analyser with a GPS module for location awareness. Significant averaging was performed, with approximately 300,000 samples per bin. The spectrum is shown below in Fig. 4.5. Also shown are indicators for the edges of the 8MHz bands in the range. Some of the bands are numbered to allow easier reference.

It is possible to estimate the power in each band from this data and, therefore, the SNR.

Firstly, the noise floor of the reciever must be found. In Fig. 4.5 band number four appears to be empty. The frequency response is flat at approximately -125dBW. Each 8MHz channel corresponds to 42 bins, thus the noise power in the channel is 10 log_{10} (42) dB higher i.e. 16.2dB higher. Therefore, the noise power in a 8MHz band is approximately

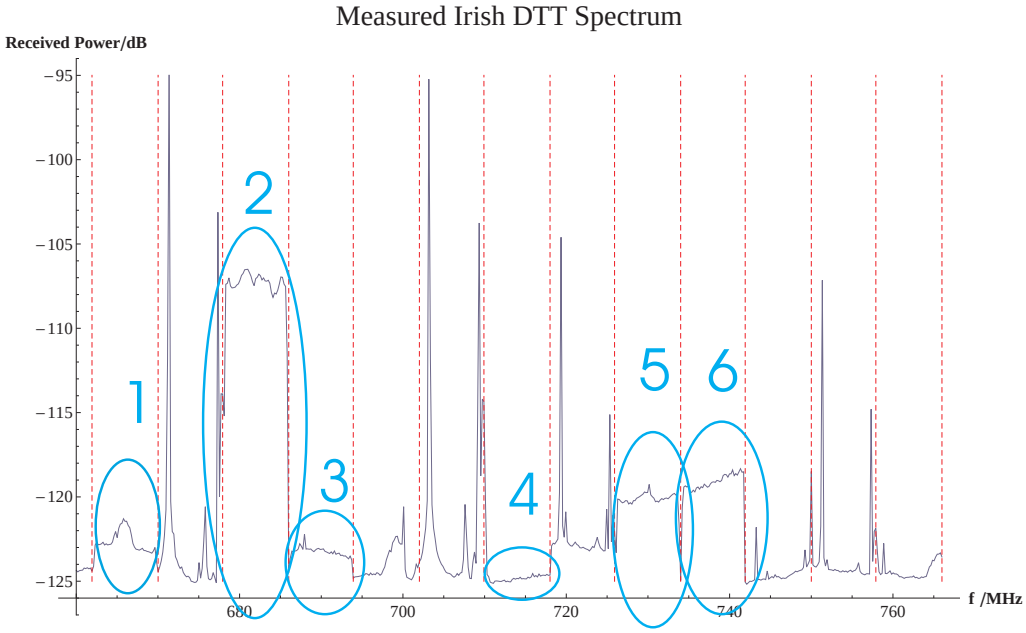

Figure 4.5: Irish DTV Spectrum Measurements

-109dBW.

The total power in a band is given by the summation of the individual powers that are in that band, as long as the signals do not create destructive interference. In this case the total power includes the noise. For the measurements considered here, the signals are OFDM signals with an 8MHz bandwidth and the power is spread across the entire 8MHz bandwidth. Note that the powers must be added in a linear scale. For example, the total power, P_{Total}, in band two is -91.2dBW .

Once the total power in the band is found then the noise power must be removed from it. Again, this must be done in a linear scale. Once this power is removed the result is the signal power in the band. Taking the ratio of the signal power to the noise power gives the SNR. For band two, the signal power, P_s, is -91.3dBW and the SNR is 17.41dB.

The results for this analysis are shown in Table 4.2.

Table 4.2: SNR of Various Bands from Measurements

Band	SNR
1	-1.45 dB
2	17.42 dB
3	-3.31 dB
4	Empty Channel
5	3.06 dB
6	4.33 dB

As can be seen, there is a range of SNR values. If a larger number of locations, including indoor locations, were used then it would be possible to determine an SNR distribution using this data.

This method, however, is of limited practical use. Measurements are only valid for that single location and time, and will become invalid if the radio is in a different location.

In addition, there is significant difficulty when attempting to find the SNR of weak signals. Reducing the noise floor to -130dBW and obtaining increased accuracy for narrowband signals would require a tenfold increase in the resolution. This would enable signal analysis at a lower SNR, though only if the signals do not cover the entire bandwidth. For bands where signals cover the entire bandwidth, such as DTV OFDM signals, significant averaging would have to be used. For example, for a signal with an SNR of -15dB, the difference between it and a noise floor at -125dB_W would be 0.135dB, a difference that is, in practice, very difficult to accurately detect. The difference between signals with an SNR of -21dB and this noise floor is only 0.035dB. Substantial averaging would have to be performed to be able to accurately decide on the occupancy and the SNR of such signals. If this process was repeated in a number of different locations and times, an SNR PDF could be created, but the system would still find low SNR signals difficult to categorise.

4.3.3 In-Line Measuring

As discussed above, the previous methods of estimating the distribution have some difficulties associated with them. In addition to those already mentioned, one major problem is that the methods must be performed in advance. If the conditions change, then the distribution changes also. Significant changes could render the distribution useless, or even detrimental to performance.

If the measurement was performed by the system as part of its spectrum sensing, then the distribution would match more closely to that of reality. In addition, as the conditions changed, the system would be able to learn and adapt the SNR distribution used. This would allow optimisation to occur during operation, a very desirable feature in CR applications.

Most detector outputs depend on the SNR of the signal, though often the relationship is complicated. In the case of cyclostationary feature detectors, for example, the magnitude of the peak depends on the cyclic correlation of the signal, in addition to its power [58]. However, energy detectors have a linear response to SNR and are ideal for this purpose. Indeed, using the energy detector output, it is a relatively simple task to estimate the SNR. Whilst this SNR estimate will not be accurate enough to be used directly [96], it is sufficient to generate an estimate of the SNR distribution.

The energy detector output, Y, has a mean μ which is given by $N + N \times SNR$ as in (3.11). Thus, the estimate of the SNR, $\hat{\gamma}$, of a signal is given by:

$$\hat{\gamma} = \frac{Y}{N} - 1. \tag{4.12}$$

There is no concern about receiving negative values for $\hat{\gamma}$. For $Y \leq N$ to occur would require $P_{fa} \geq 50\%$, clearly a very poorly designed detector and, thus, it is safely assumed

here that this will not occur.

This method of estimation would appear to be the most promising. Though a number of detection attempts are required to "train" the CR to generate the correct distribution, it has the ability to learn and adapt. In addition, it can find the SNR of weak signals reliably as the energy detector must be able to detect them. Finally, this method does not require estimation of channel loss or any of the other parameters that are required for the analytical method.

There are, however, practical issues which reduce the accuracy of this estimation technique. The most important factor in the practicality of the method is the required number of observations before the distribution is sufficiently accurate. How many times does the detector need to be run before the distribution is appropriately reliable? This is now investigated in simulations and in practical work.

4.4 Analysis of the In-Line Measurement System

There exist some non-idealities in the SNR estimation method of equation (4.12). Noise uncertainty and the effects of the threshold and occupancy on the system generally renders the estimates less accurate.

4.4.1 Effects of Noise Uncertainty on SNR Estimation

The estimation of the SNR distribution requires the use of an energy detector to generate an estimate of the SNR. The energy detector and, thus, the SNR estimate, is susceptible to noise uncertainty. Noise uncertainty causes a spreading in the output PDF of the energy detector (i.e. the variance increases).

Assuming the noise uncertainty is unbiased, then the energy detector output remains unbiased also. Whilst the mean of the energy detector's output does not change, the variance increases. This means that extra samples will be required to generate an accurate result, compared to the case where no noise uncertainty is present. For low values of noise uncertainty, there will be no significant error in the distribution. As the noise uncertainty increases, however, the effects of the threshold and occupancy will also increase. This will lead to errors in the estimated distribution, even when the number of estimates is very large. In addition, if the noise uncertainty has a bias, then the estimated distribution will be incorrect, even when the number of estimates is very large. The performance of the energy detector in-line measurement system under noise uncertainty is shown later in Fig. 4.8.

4.4.2 Effects of Threshold and Occupancy on SNR Estimation

The in-line SNR estimation technique will give the correct distribution when the channel is occupied and there is no threshold. In real applications, however, the channel may not be occupied and a threshold will be present.

To see the impact of these factors, consider the energy detector output PDF in section 3.3. The unoccupied channel (H_0) has a probability of being declared occupied of P_{fa}. This corresponds to the section of the H_0 PDF that is above the threshold, λ. In addition, the SNR estimate will be low due to the fact that the more probable values for the energy detector output are closer to the threshold. Thus, there is a low, incorrect, value for the SNR with a probability of occurring of $P_{fa} \times (1 - \theta)$, where θ, $0 \leq \theta \leq 1$, denotes the occupancy of the channel. For a channel with 50% occupancy and $P_{fa} = 10\%$, approximately 5% of the detection attempts will result in incorrectly designated unoccupied

channels with a low SNR estimate.

Similarly, for the occupied channel (H_1), there is a probability of being declared unoccupied of P_{md}. In this case, the lower values of the output will be removed, corresponding to the section of the H_1 PDF that is below the threshold, λ. This censoring of the occupied channel results in the lower SNR values being removed. Thus, there is a low value for the SNR which is ignored, with a probability of occurrance of $P_{md} \times \theta$. For a channel with 50% occupancy and $P_{md} = 10\%$, approximately 5% of the detection attempts will result in incorrectly designated occupied channels with a low SNR estimate being ignored.

It can be seen that the two effects will counteract each other somewhat. The new distribution from which the SNR estimate is generated, compared to the H_1 distribution, from which it is assumed to be generated, is shown in Fig 4.6.

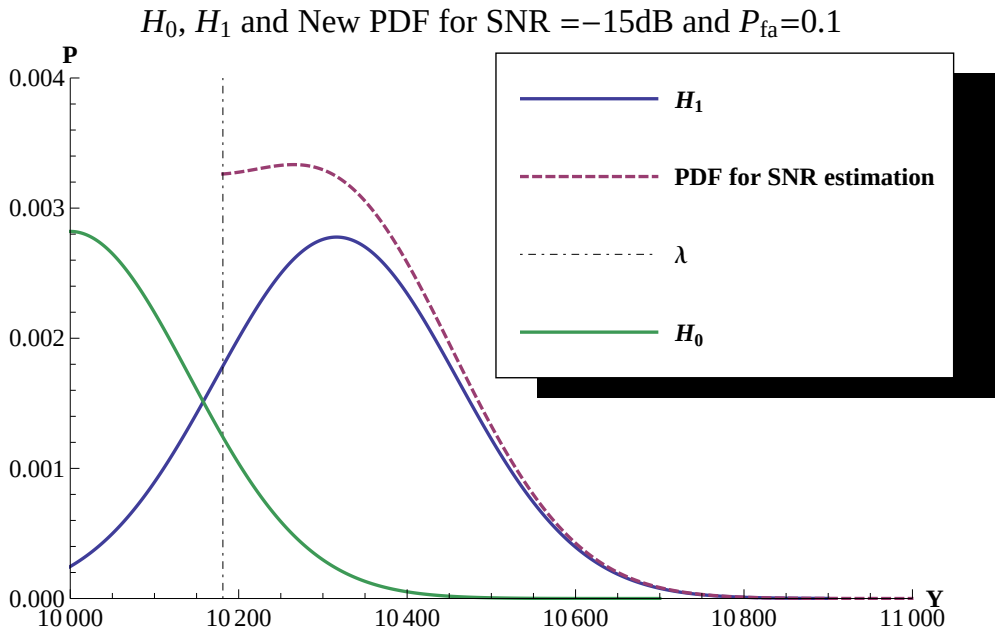

Figure 4.6: PDF from which the SNR Distribution is Actually Based vs H_1 Distribution

It might be possible to reduce the impact of the threshold and occupancy if the occupancy is known. If a sufficiently large number of estimates is taken, then, by removing some of the lowest valued results, which would correspond to the unoccupied channels incorrectly designated, then some of the bias would be removed.

The PDF at each point could also be scaled by a term to counteract the bias. This term would depend on the values of θ and P_{fa} and the value of the SNR that corresponds to that point on the PDF. To test this method, Monte-Carlo simulations and practical tests were performed using Mathematica® 8.0.1.0 and the IRIS system, respectively.

4.4.3 Simulations

The simulations were performed by generating energy detector outputs based on an SNR distribution. The SNR PDF used here is the distribution based on the Irish DTV network derived in section 4.3.1. Initially, a random variable from this distribution is generated for the SNR and then, using this SNR value, the output for the energy detector is generated. Using (4.12), the estimate of the SNR is generated. This process is repeated N_{obs} times. Once all N_{obs} observations have been collected, the KDE method is performed and the estimated distribution, $\hat{f}$, generated.

In Fig. 4.7 sample distributions for N_{obs} = 10, 100 and 1000 are shown, as well as the actual distribution. This is for the case where the threshold and occupancy have to be accounted for. The bias in the SNR distribution can be seen.

The $MISE$ is then calculated for f and $\hat{f}$. The process is repeated 5000 times and the average value for the $MISE$ found. The result for the DTV distribution is shown in Fig. 4.8. In addition, the results under noise uncertainty and when the effects of the threshold and occupancy are included are shown.

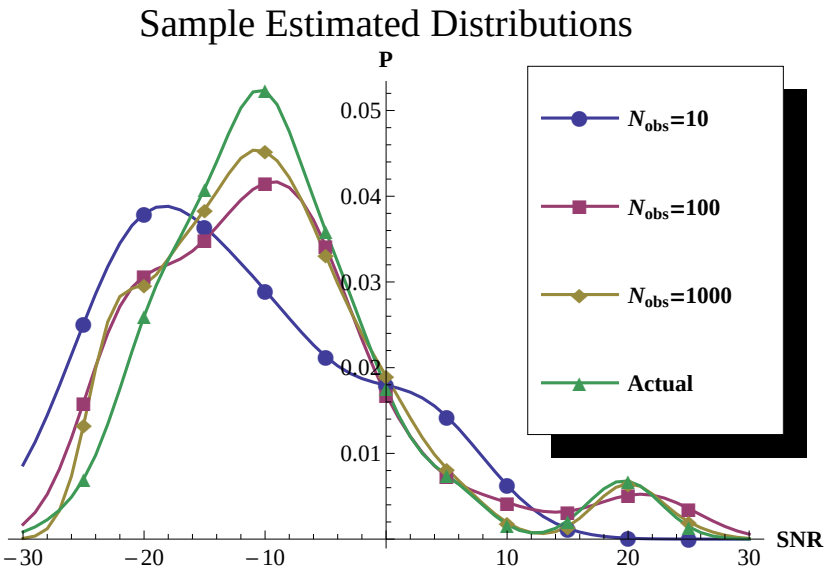

Figure 4.7: Sample $\hat{f}$ for Various Observation Lengths, N_{obs}

There is a log-log relationship between the number of observations and the $MISE$. To reduce the $MISE$ by a factor of ten then the number of samples required increases also by approximately the same factor. This has only been shown for the distribution used here though, and other distributions will likely have different characteristics, especially multimodal distributions. It is expected that other distributions would have similar responses, but this cannot be proven for all distributions. Thus, to guarantee performance, the $MISC$ is the recommended test for the 'fit' of the distribution.

The presence of noise uncertainty reduces the effectiveness of this system. There is a greater error compared to the case where no noise uncertainty is present. In addition, greater numbers of samples do not reduce the $MISE$ by as much as the case where no noise uncertainty is present. The noise uncertainty is quite large in this case though, noise

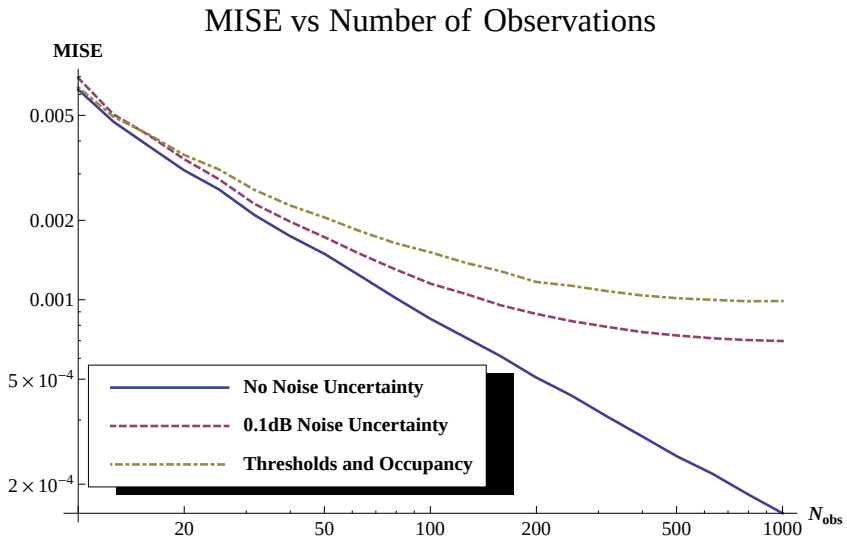

Figure 4.8: $MISE$ for Differing Numbers of Observations,N_{obs}

uncertainty of 0.1dB would be too large for the energy detector to reliably detect signals. Lower levels of noise uncertainty would cause less deterioration in the performance of the system.

The degradation due to the threshold and occupancy, however, cannot be ignored. There is a greater error compared to the case where noise uncertainty is present. In addition, greater numbers of samples do not reduce the $MISE$ by as much as the case where noise uncertainty is present. This impairment of performance will be present in the detector under all conditions; thus, it must be included in all calculations.

The $MISE$ values shown are the average $MISE$ values for that number of observations. In some cases the $MISE$ will be larger than this. To provide a measure of confidence for the estimated distribution, the $MISC$ can be used.

For the DTV distribution and $k = 2$, i.e. doubling the number of samples each check, the $MISC$ is calculated starting with N_{obs} = 20. This process is repeated 5000 times, and the $MISC$ is calculated at each step. The average $MISC$ is shown in Fig. 4.9. These simulations were performed on the boole cluster, NOTE

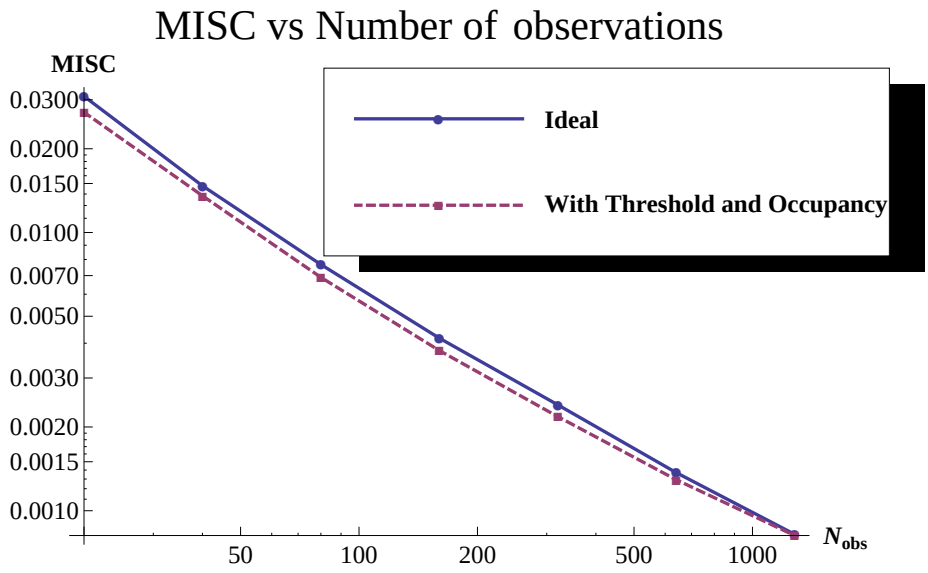

Figure 4.9: $MISC$ for Various Observation Lengths

Finally, for a practical system, a target $MISC$ value will be specified. Once the system has a smaller $MISC$ than the target $MISC$ value, $\hat{f}$ can be considered stable and sufficiently accurate.

Shown in Fig. 4.10 is a plot illustrating the probability of having met the target $MISC$ by that number of observations for various target $MISC$ values. Obviously, the smaller the $MISC$ allowed, the longer the system takes to stabilise on a distribution.

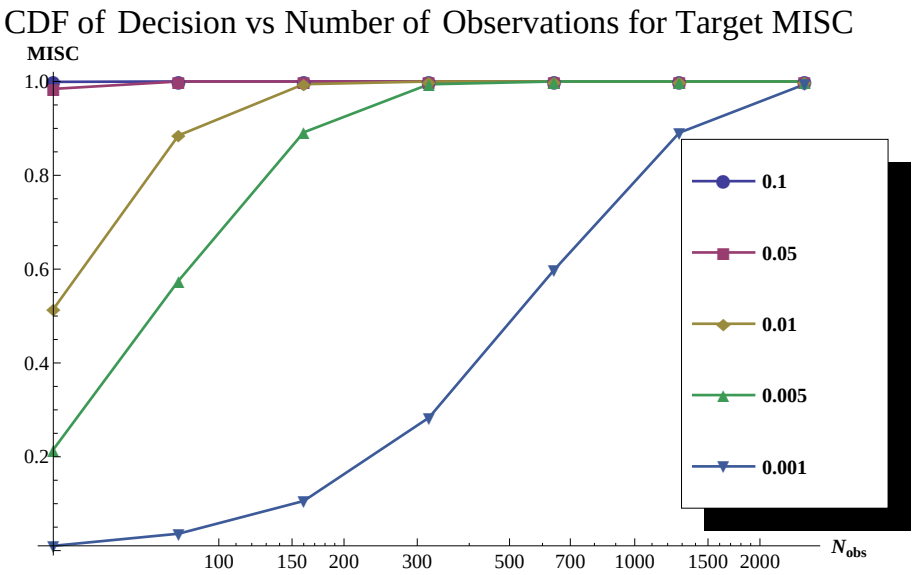

Figure 4.10: Probability of Reaching Target $MISC$ vs Number of Observations

4.4.4 IRIS Implementation and Testing

To test this method in reality, a radio was designed on the IRIS system that would estimate the SNR PDF. A simple two radio system was used for this test, with a second radio generating a random signal. The designs for both radios are shown in Fig. 4.11.

The transmitting radio, Tx, generates random data for an OFDM modulator. This is then sent to the USRP front-end and transmitted over the channel. The power of the signal is controlled by a gain module. The effective SNR at the receiver could also have been controlled by modifying the bandwidth of the signal, though this was not done as the gain module provided sufficient flexibility.

The receiving radio, Rx, receives samples from a USRP front-end at a sample rate of 1MSps and passes them into an energy detector component. The results from the energy

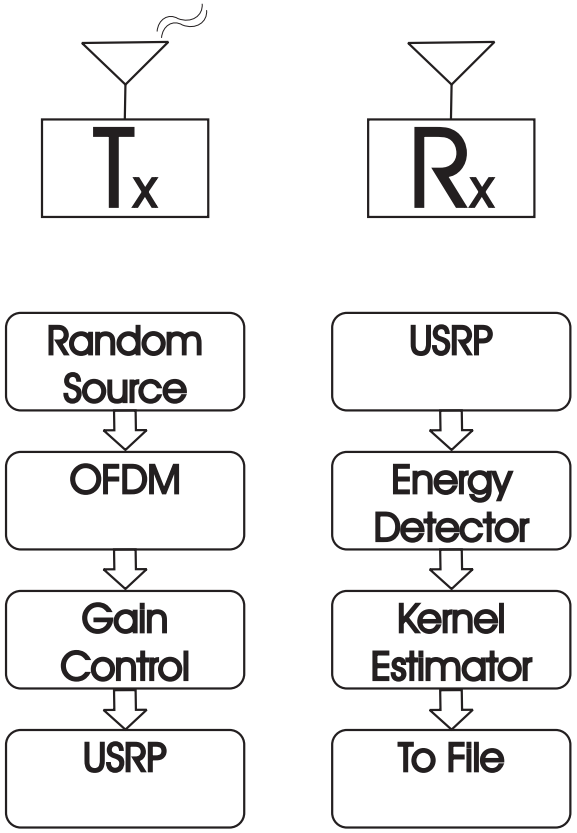

Figure 4.11: IRIS Test Radio Architecture

detector are fed into the KDE component which then performs one of the tests and the results are written to a file. Both radios were operated at a frequency of 5.008GHz and a photo of the equipment is shown in Fig. 4.12.

The time taken for Rx to generate an estimate was found to be approximately 0.2 seconds (209k samples at 1MSps). Tx changed its gain approximately every 0.3 seconds. This means that no two sensing periods had the same SNR. In a real system Rx would not be scanning the same band repeatedly, thus larger periods of time would pass between sensing attempts than in this test. In addition, multiple bands would be scanned when generating the SNR distribution. Thus, it is unlikely that the same SNR would be found

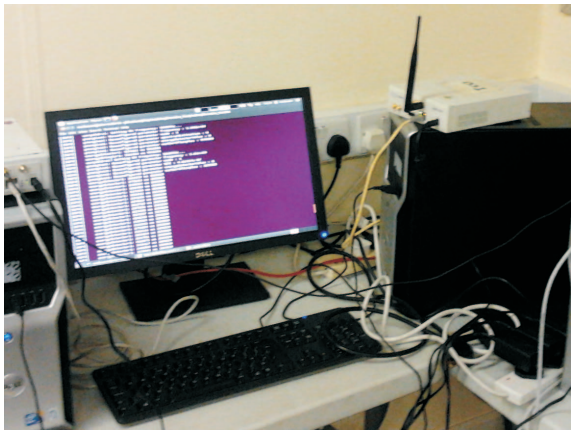

Figure 4.12: IRIS Setup in Trinity College Dublin

for two sequential estimates. Hence, the constantly changing SNR is justified, as it models a more realistic environment. It may be possible that the transmitter and receiver are static over long periods of time, thus the SNR in that band would not change significantly with time. In this case, however, the multiple bands being scanned still result in a range of SNR values at the receiver.

The effective SNR, $\gamma_{effective}$, at Rx is given by:

$$\gamma_{effective} = \frac{P_s}{P_n}. \tag{4.13}$$

The Rx USRP has a bandwidth of 1MHz in this experiment. This does not change and, therefore, the noise power remains constant.

The USRP front end using a XCVR 4250 daughterboard displayed a drift in gain with time. As the device began receiving samples the gain decreased. This is due to the rise in temperature in the device changing the gain. Thus, as the device continued to operate, the gain would slowly decrease. Once a certain threshold was reached the gain would shift suddenly as automatic gain controls attempted to compensate. At this point in all the tests

the system was shut down and allowed to cool.

The first test performed measured the average $MISE$ for a range of N_{obs} values. In this test the gain began relatively high and then decreased slowly over time, with a 10% change in gain over approximately an hour. As the decrease was slow and of relatively low magnitude the estimation did not suffer significantly. The average distributions for the N_{obs}=10, 50 and 100 estimates are shown in Fig. 4.13. The tests were performed sequentially starting with the N_{obs}=10, then N_{obs}=50 and, finally, N_{obs}=100. Indeed, the effect of the gain drift could be seen as analogous to noise uncertainty. The scaling being performed is incorrect, thus the estimates will suffer from the same effects as noise uncertainty. A sample distribution from each of the tests is shown in Fig 4.14.

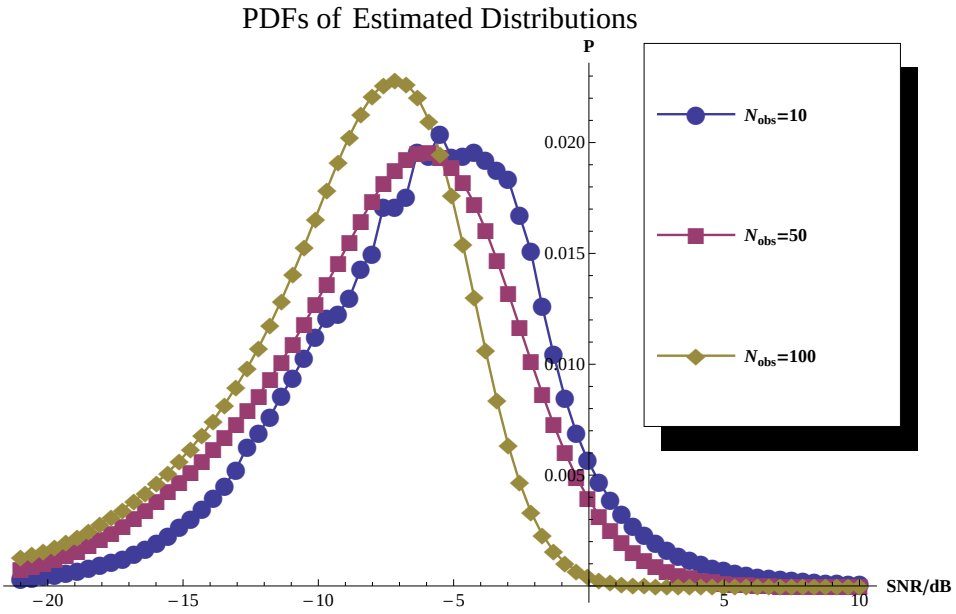

Figure 4.13: Change in the SNR Distribution Between the N_{obs}=10, 50 and 100 Estimates runs

Unlike the simulation, the underlying distribution for the IRIS test is not known. The

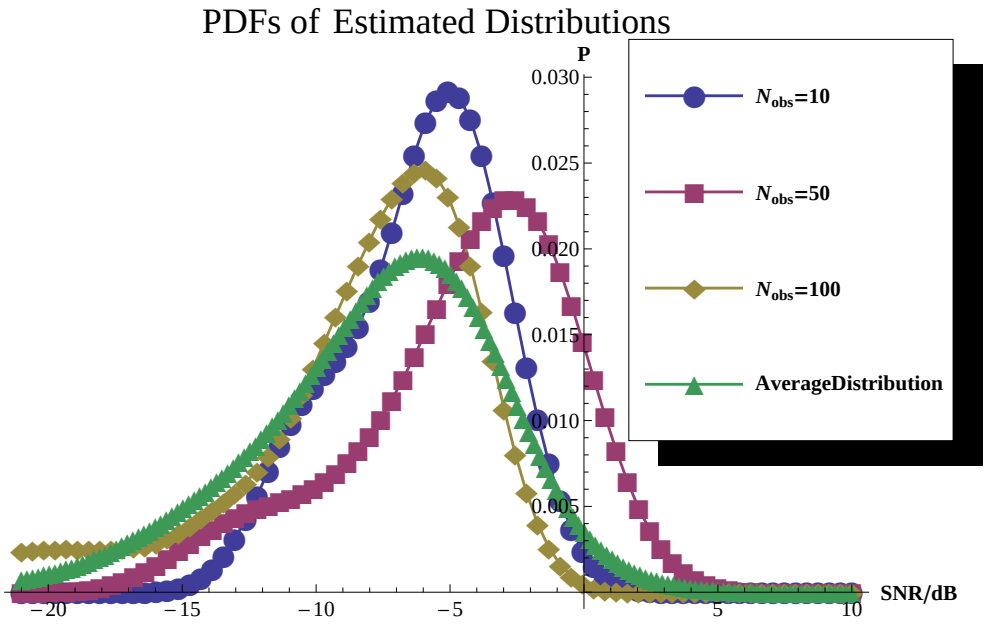

Figure 4.14: SNR Distributions for N_{obs}=10, 50 and 100 Estimates.

underlying distribution is estimated from an average of all the distributions found. Each estimated distribution is compared to a distribution that is the average of all the distributions with the same number of estimates (i.e. the N_{obs} =10 estimated distributions are compared to the average of all the N_{obs} =10 estimated distributions). Thus, any bias introduced by the thresholds would also be in the distribution with which each distribution was compared. Therefore, the correct simulation condition with which to compare this test is the noise uncertainty case. The $MISE$ is then generated from the estimated distributions. The $MISE$ is shown in Fig. 4.15.

The results do not match well. The $MISE$ is significantly higher than predicted. This could be caused by the noise uncertainty and by the gain drift of the system. Approximately 5x samples are required by the practical system to have the same performance as

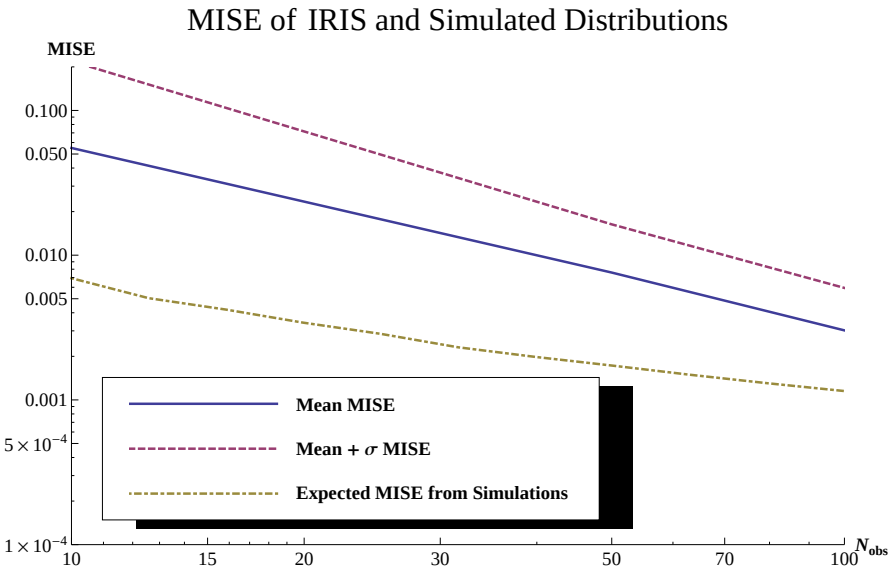

Figure 4.15: $MISE$ for IRIS Test

the simulated system.

This mismatch shows the need for constant noise power estimation in an adjacent band to reduce the effects of the gain drift. The $MISC$ test described next used noise power estimation in an adjacent band, to reduce the impact of the drift in gain.

As has already been stated in this chapter, the $MISE$ is of little practical use to a system because it requires knowledge of the distribution being estimated. The $MISC$, however, as defined in (4.9), can be used to provide a measure of confidence.

As the test is sequential (each step uses the previous step to compare the current SNR estimate to) variance in the SNR will lead to difficulties. Thus, an extra controller was added to the system. Every fifty energy detector results the system would change to an adjacent (empty) band and recalibrate the gain to compensate for the thermal effects. The

system would then return to the original band and resume the test.

Finally, the CDF of the $MISC$ simulation is shown. This is the probability that the system will have reached an estimate by the n^{th} estimate for the $MISC$ values chosen. This is shown in Fig. 4.16.

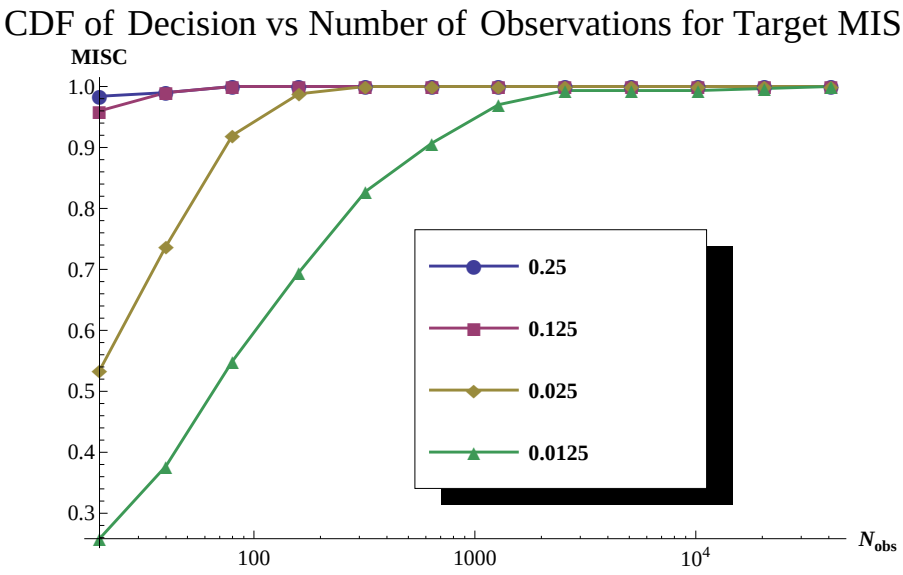

Figure 4.16: $MISC$ CDF for Decision on Distribution for IRIS Test

Note that the $MISC$ target values are not very low and are easily met by the system in most cases. However, the systems that will use the distribution are robust to inaccurate distributions, as will be shown later in Chapter 5 and Chapter 6. Therefore, the $MISC$ target values do not need to be very low, thereby allowing a small number of estimates to be used when estimating the SNR distribution.

4.5 Conclusion

In this chapter the ability of a system to generate an accurate SNR distribution was investigated. In section 4.3, various options for SNR distribution estimation were examined. The theoretical method suffered from a large degree of uncertainty in the parameters to be used and was rejected as a viable option. Likewise, the experimentally in advance method was also rejected though, in this case, the main issue was the fact that conditions could change between having learned the distribution and the usage of that knowledge.

The experimentally in-line method was selected as the most promising. The method was shown to have little cost, as the individual SNR estimates are readily available.

Using KDE with this method allows relatively quick and accurate estimates of the SNR distribution to be found. In section 4.2, two methods for assessing the accuracy of a distribution were presented. The first, $MISE$, is a well-studied method for checking the accuracy of an estimate when the original distribution is available. When the original distribution is unavailable, $MISC$ allows the system to check that the distribution has stabilised, a good indication that the distribution is accurate. This $MISC$ was developed for the SNR estimation scheme here, though it is likely a similar technique exists in other work under a different name.

There are some non-idealities in the in-line estimation method. Noise uncertainty can reduce the accuracy of the system. In addition, the threshold of the detector and occupancy of the channel can introduce a bias in the estimates. Although it may be possible to reduce this bias, this is outside the scope of this work.

The method was tested both in simulations and practically using the IRIS system. A drift in the gain of the USRP was noted which reduced the accuracy of the system. This was corrected by the system taking very regular estimates of the noise power in an adjacent

band. The IRIS results are not the same as the predicted ones from simulation. However, the estimated distributions are still sufficiently close to allow the detector architectures in Chapters 5 and 6 to work.

Finally, it should be stated that this method is not intended to generate an SNR distribution to be used to guarantee safety to primary users. The SNR distribution should only be used to increase performance in a way that does not change the probabilities of interference with the licensed users of the band. In this work the SNR distribution is only used to decrease the sensing time required. Thus, even if the estimate is substantially incorrect, the primary user will not suffer, rather the CR alone will have reduced performance.

5

Coarse Sorting Fine Detector

5.1 Introduction

5.1.1 Coarse-Sorting Fine Detector Architecture

When a cognitive radio attempts to find a free channel, it will not check a single channel only. Instead, several candidate channels will be investigated. If the CR is attempting to find the first free channel available, then it will stop once it has found a channel that

it declares to be free. Under most conditions, aside from heavily correlated channels, a random search is the optimum, when no extra information is available [79].

However, if extra information was available, this would no longer be the case. If a quick scan was performed on all the channels, then extra information would be available to the CR. This would allow the CR to make a better decision on which of the channels to intensively scan first. This could be achieved by first scanning all the channels with an energy detector with a small number of samples per detection attempt, this is the coarse detection phase. Then the channels are sorted w.r.t. the outputs of the coarse scan. The channels would then be scanned by an accurate detector, this is known as the fine detector. Such a detector has been known as a coarse-fine detector. However, to differentiate it from other detectors that use coarse and fine scans but without any sorting, such as will be seen in Chapter 6, it is denoted here as a Coarse-Sorting Fine Detector (CSFD). The flow diagram associated with a CSFD is shown in Fig. 5.1. Note that "next lowest" is the lowest chanel on the first iteration. Also included is the flow diagram for a Naive detector, where there is no coarse detection and, therefore, no sorting. An important question is, how coarse should the coarse detection segment be to ensure efficient and effective operation?

Consider the environment in Fig. 1.1. Some of the signals are strong and at a high SNR. These signals would be easily detected by a quick (coarse) scan. Other signals are weak and have very low SNRs. These signals would not be detected reliably with a small number of samples. However, the coarse detector is not intended to detect the signals reliably. The coarse detector is attempting to provide an estimate that can be used to sort the channels. As the number of samples in the coarse detector increases, the effectiveness of the sorting will increase also.

To find the optimum point, corresponding to the minimum required number of coarse

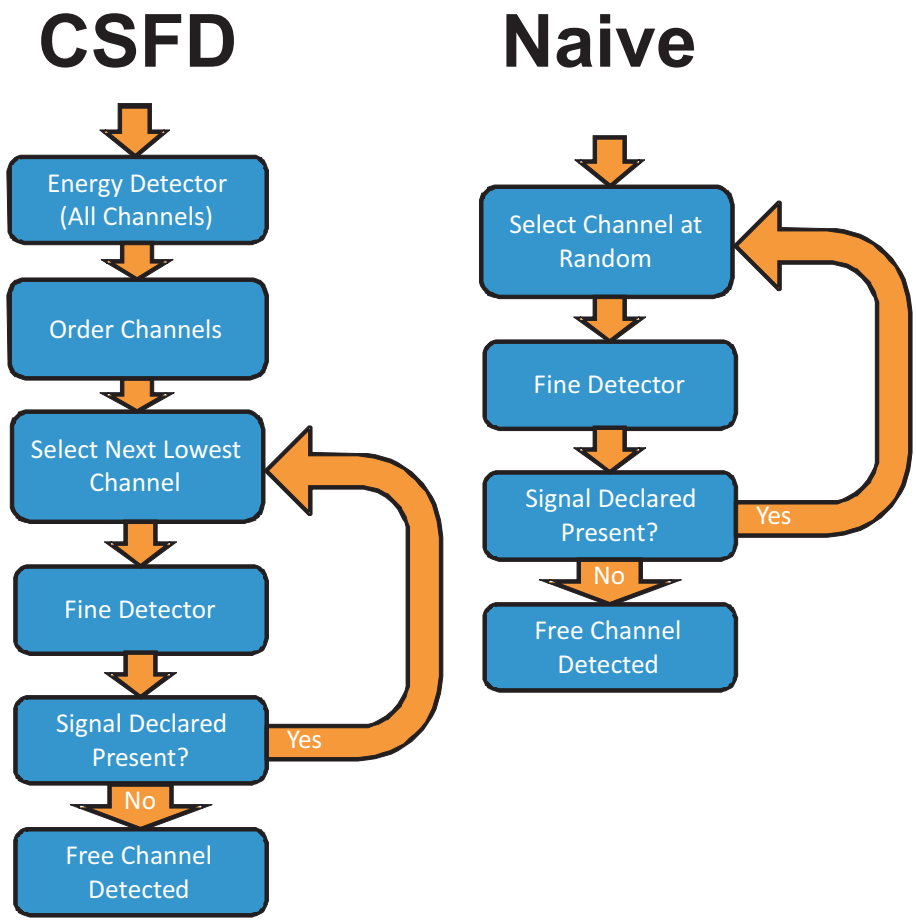

Figure 5.1: CSFD Flow Diagram Compared to Naive Detector

samples where the gain of taking extra samples is no longer worth the associated cost, is difficult analytically. The SNR distribution may be user-defined and expressed as a sampled distribution. Thus, in general, an analytical solution is highly likely to be intractable. Monte Carlo simulations could be performed to find the optimum point. In practice, these simulations would most likely be accurate but costly in terms of computations and time. Instead, a model of the process was generated that allows the accurate prediction of system performance in significantly less time than Monte-Carlo simulations running on the same hardware.

5.1.2 Modelling the System

First, an expression for the PDF of the output of the energy detector is found. This PDF accurately describes the output of the energy detector over different numbers of samples and the prevailing SNR distribution, if a signal is present. In addition, the effects of noise uncertainty on the distribution of the energy detector output are modelled.

Using order statistics [110], the outputs from the energy detector are sorted. It is assumed that the SNR distributions are independent and identically distributed (i.i.d.). This assumption reduces the computational complexity of the sorting process dramatically, without impacting the accuracy of the estimates of the sorted outputs. By comparing the sorted outputs, the probability that a signal is present in the channel analysed in the n^{th} detection attempt is calculated. The reduction in the accuracy of the model using this method for a channel with a time varying occupancy is also found, allowing more realistic modelling.

Once the probabilities of occupancy have been found for all the channels, Markov Chains can be used to predict the performance of the system. Using the Fine Detector's Receiver Operating Characteristic (ROC) the performance can be estimated, even in the presence of fading. The fine detector could have a different ROC to the coarse detector if a different detector is used for the coarse and fine detectors. In this case an energy detector is used for both and there is only one ROC to consider. The Markov chain can be used to find the average number of samples required to find a free channel (and, thus, the speed gain relative to a Naive detector), the variance in the number of samples required to find a free channel, the probability of generating interference and the probability of not finding a free channel.

In this chapter the various components of the model are described and the accuracy of each section is discussed. Finally, the overall accuracy of the model is shown by comparisons

to large scale Monte-Carlo simulations.

5.2 Energy Detector Output

For an energy detector based coarse detector the output, Y, as derived in [54], is chi-square distributed when there is no signal present (H_0) and non-central chi-square distributed when a signal is present (H_1) or,

$$Y \sim \begin{cases} \chi^2_{N_{coarse}} & H_0 \\ \chi^2_{N_{coarse}}(\gamma) & H_1 \end{cases}, \tag{5.1}$$

where N_{coarse} is the number of degrees of freedom (equal to the number of samples used by the coarse detector) and γ is the non-centrality parameter, given by the product of the SNR and N_{coarse}.

The distribution of the output, when no signal is present, depends only on the number of samples used. In Fig. 5.2, the output PDF is shown for various values of N_{coarse}.

When a signal is present the output depends on both the number of samples and the SNR. In Fig. 5.3, the output PDFs are shown for various values of SNR.

The distribution of the output of the energy detector, Y, depends on the distribution of the SNR. If the SNR distribution is known, then the new distribution, $f_{Y,SNR}$, can be found by averaging over the SNR via:

$$f_{Y,SNR}(x) = \int_0^\infty f_{SNR}(\gamma)\chi^2_{N_{coarse}}(\gamma, x)\mathrm{d}\gamma, \tag{5.2}$$

where $\chi^2_{N_{coarse}}(\gamma, x)$ is the value of the non-central chi squared distribution at x, for a non-centrality parameter γ and N_{coarse} degrees of freedom and $f_{SNR}(\gamma)$ is the value of

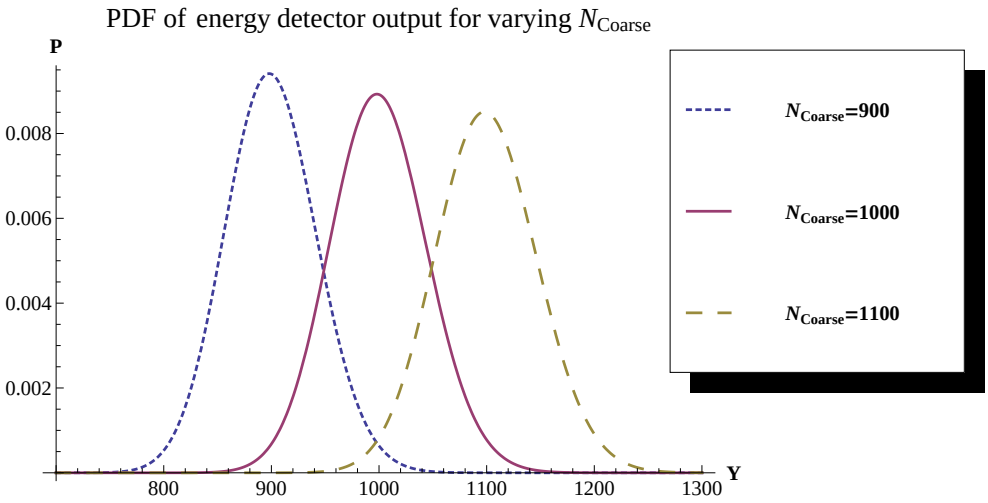

Figure 5.2: PDF of Energy Detector Output When No Signal is Present

the SNR PDF for an SNR of γ.

Since, in practice, most distributions will be derived from observations, as in [106], expressing $f_{Y,SNR}$ in terms of a sampled SNR distribution is usually more appropriate. In addition, since the dB scale contains more information for a uniformly sampled SNR distribution, the range is changed to dB, thus (5.2) becomes,

$$f_{Y,SNR}[x] = \Delta \sum_{\gamma=-\infty}^{\infty} f_{SNR}[\gamma]\chi^2_{N_{coarse}}[\gamma, x], \tag{5.3}$$

where Δ is the step size of the summation and the range of the summation reflects the appropriate bounds. In the example used in this work, the SNR distribution is in the range -21dB to 10dB and so the summation is performed over this range. Note that the operation is not, strictly speaking, a summation; the step size is not usually equal to unity. The summation notation is used for expository simplicity.

The SNR distribution derived in the previous chapter is used for demonstrative purposes

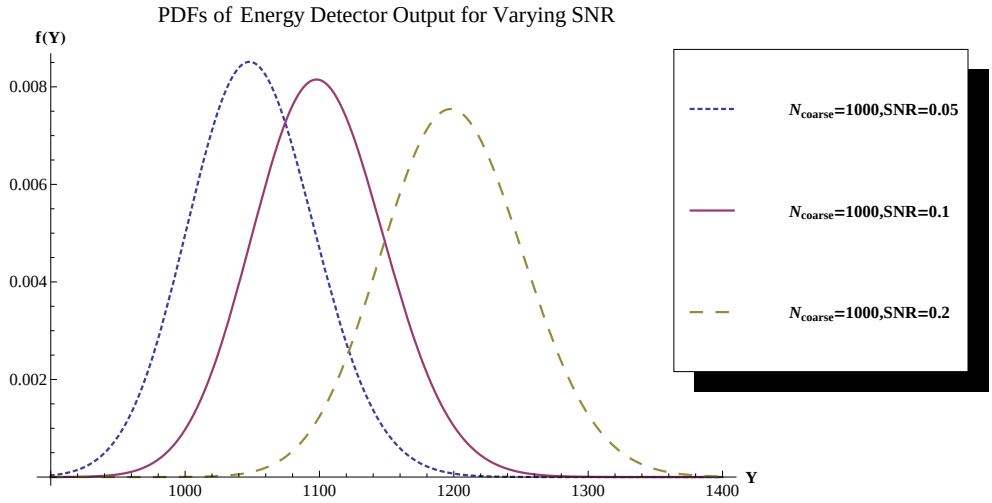

Figure 5.3: PDFs of Energy Detector Output When a Signal is Present

throughout this chapter. For this example distribution, the PDF of the output of the energy detector is illustrated in Fig. 5.4.

This distribution, $f_{Y,SNR}[x]$, is then the distribution of the output of the energy detector when a signal is present, without the detector having any knowledge of the SNR other than its distribution. The small peak at roughly 11000 is the peak seen in the sample SNR distribution at approximately 20dB

5.3 Energy Detector Outputs under Noise Uncertainty

Thus far in the analysis it has been assumed that the system has accurate knowledge of the noise variance, σ_n^2. This variance is used to scale the input of the energy detection to give a unit variance input when no signal is present. However, if the estimate of the variance, $\hat{\sigma}_n^2$, is incorrect, then the energy detector input will be scaled incorrectly.

ρ is defined as the noise-uncertainty, or the relative inaccuracy of the estimate, and is

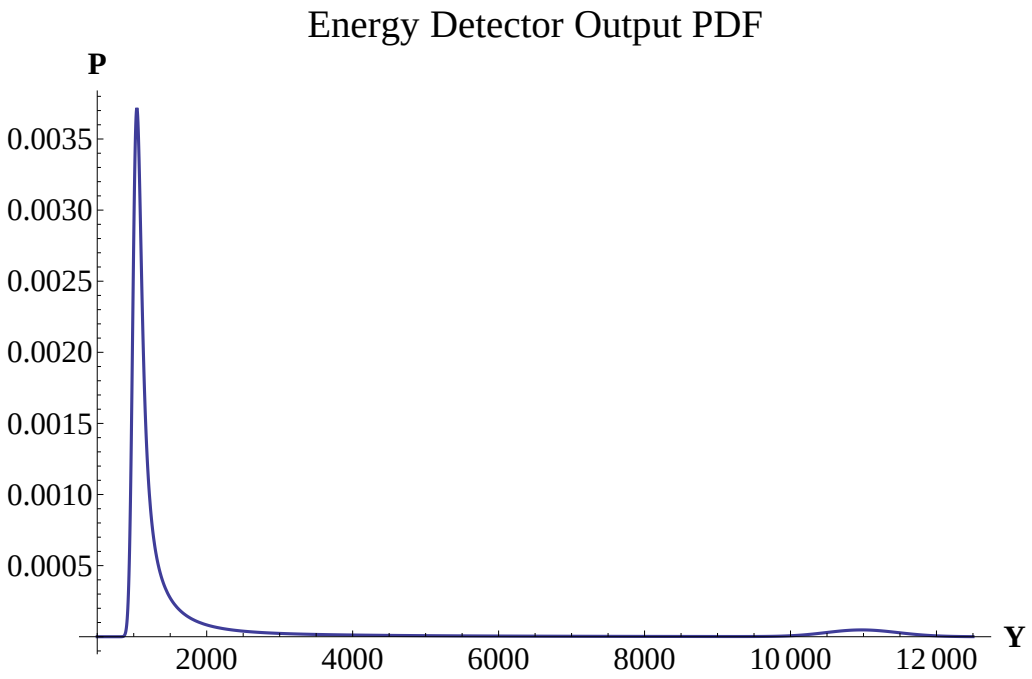

Figure 5.4: PDFs of Energy Detector Output When a Signal is Present With SNR Distribution of Fig. 4.4

given by [42]:

$$\rho = \frac{\hat{\sigma}_n^2}{\sigma_n^2}. \tag{5.4}$$

5.3.1 Noise-Uncertainty for Unoccupied Channels

For a noise-only channel, the distribution of the input to the energy detector, after normalisation, has variance ρ:

$$v(n) \sim \mathcal{N}(0, \rho). \tag{5.5}$$

The energy detector output:

$$Y_n = \sum_{n=1}^{N_{coarse}} v^2(n), \tag{5.6}$$

is then distributed according to:

$$Y_{n,\rho} \sim \sum_{n=1}^{N_{coarse}} \mathcal{N}(0,\rho)^2, \tag{5.7}$$

which can be written as:

$$Y_{n,\rho} \sim \rho \sum_{n=1}^{N_{coarse}} \mathcal{N}(0,1)^2. \tag{5.8}$$

The sum of the squares of N_{coarse} zero mean, unit variance, Normally distributed random variables is χ^2 distributed with N_{coarse} degrees of freedom.

The PDF of a noise-only energy detector output with uncertainty ρ, $f_{Y_{n,\rho}}[x]$, is, therefore,

$$f_{Y_{n,\rho}} \sim \rho \chi^2_{N_{coarse}}, \tag{5.9}$$

$$f_{Y_{n,\rho}}[x] = \frac{1}{\rho} \chi^2_{N_{coarse}}[\frac{x}{\rho}]. \tag{5.10}$$

To generate the PDF of the energy detector output, f_N, for a distribution of ρ, denoted f_ρ, these PDFs must be averaged over f_ρ. In this work a discrete distribution for f_ρ is considered. In Fig. 5.5, the output PDF is shown for various values of ρ.

Consequently, f_N becomes,

$$f_N[x] = \Delta \sum_{\rho=\rho_{min}}^{\rho_{max}} f_\rho[\rho] f_{Y_{n,\rho}}[x], \tag{5.11}$$

where, ρ_{max} and ρ_{min} are the maximum and minimum values of ρ, respectively, and Δ is

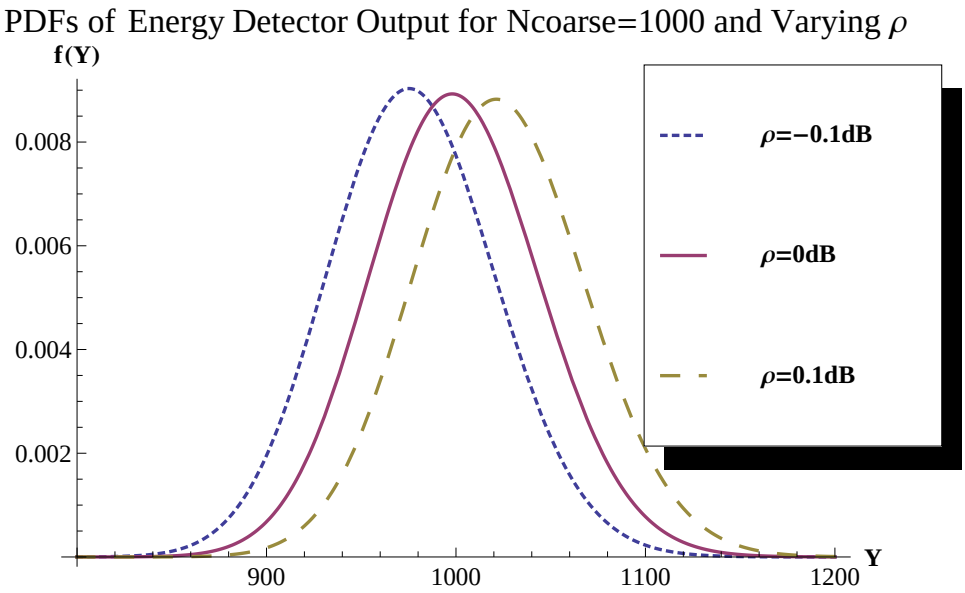

Figure 5.5: PDF of Energy Detector Output With Noise Uncertainty and No Signal Present

the step-size of the summation.

The distribution of $f_N[x]$ in (5.11) can then be used for the unoccupied case, $H(0)$, in (5.1) .

5.3.2 Noise-Uncertainty for Occupied Channels

For an occupied channel the input to the energy detector, $x(n)$, after normalisation, is the sum of two zero mean Gaussian variables, namely a noise variable with variance ρ and a signal of variance $\frac{\gamma}{\rho}$. Therefore, $x(n)$ is the sum of two Gaussian variables and has the distribution:

$$x_s \sim \mathcal{N}(0, \rho + \frac{\gamma}{\rho}). \tag{5.12}$$

Since $\rho \approx 1$ [42] and, for SNRs where noise uncertainty has the greatest effect, $\gamma \ll 1$, then (5.12) can be well approximated by:

$$x_s \sim \mathcal{N}(0, \rho + \gamma). \tag{5.13}$$

This assumption was compared with the exact distribution for $x(n)$. For $\rho = \pm 0.1dB$ and $\gamma = -10dB$ the percentage difference is approximately 1.04%. Thus, the approximation in (5.13) is deemed to be sufficiently accurate here.

When a signal is present the energy detector output is well approximated by:

$$Y_{s,\rho,\gamma} \sim \sum_{n=1}^{N_{coarse}} \mid \mathcal{N}(0, \rho + SNR) \mid^2 . \tag{5.14}$$

where the the absolute value operator is required for complex valued signals.

Therefore, the PDF of the energy detector output for a signal-and-noise channel with uncertainty ρ and SNR of γ, $f_{Y_{s,\rho,\gamma}}[x]$, is,

$$f_{Y_{s,\rho,\gamma}}[x] = \frac{1}{\rho + \gamma} \chi^2_{N_{coarse}}(\frac{x}{\rho + \gamma}). \tag{5.15}$$

In Fig. 5.6, the output PDF is shown for various values of ρ.

To generate the PDF of the energy detector output, $f_{Y,SNR}$, for a distribution of ρ, f_ρ, and a distribution of SNR, f_γ, the PDFs could be averaged separately:

$$f_{Y,SNR}[x] = step_{size\gamma} \sum_{\gamma=\gamma_{min}}^{\gamma_{max}} f_\gamma[\gamma] step_{size\rho} \sum_{\rho=\rho_{min}}^{\rho_{max}} f_\rho[\rho] f_{Y_{s,\rho,\gamma}}[x], \tag{5.16}$$

where $step_{size\gamma}$ and $step_{size\rho}$ are the step-size of the γ and ρ summations and γ_{max} and γ_{min} are the maximum and minimum SNR values, respectively.

This method, however, is computationally inefficient, requiring the calculation of the val-

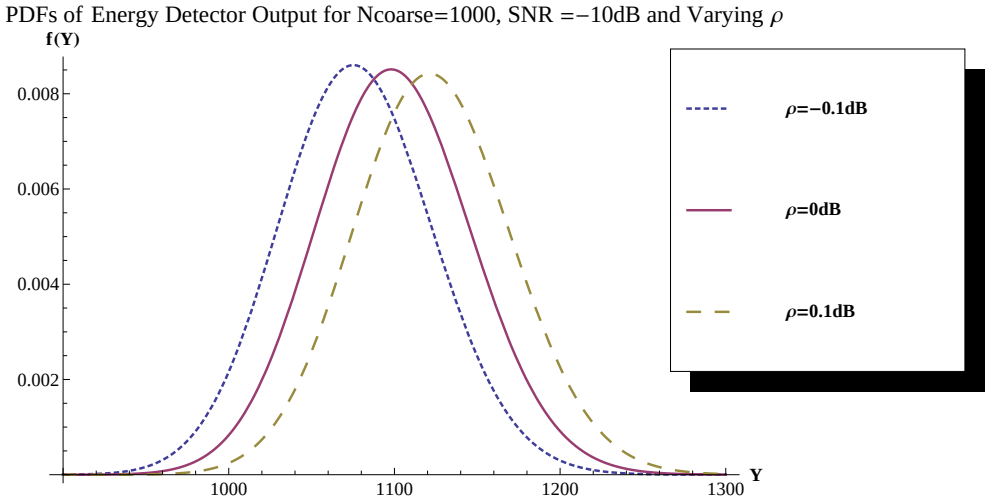

Figure 5.6: PDF of Energy Detector Output With Noise Uncertainty and a Signal Present

ues of $\chi^2_{N_{coarse}}[\rho, \gamma]$ for all values over both distributions.

Alternatively, if the PDF of $\rho + \gamma$, denoted, $f_{\rho+\gamma}$, is generated first, then the computation time can be significantly reduced. Firstly, the SNR and the noise uncertainty are converted from their dB representation to their linear form, as it is in this scale that they are added. This can be done using the cumulative distribution function and changing the scale [111].

Then the PDF of the sum of these two variables is the convolution of their respective PDFs [112], i.e.:

$$f_{\rho+\gamma}[x] = \sum_{n=1}^{N} f_{\gamma}[x] f_{\rho}[n-x]. \tag{5.17}$$

This relatively quick convolution reduces the number of times that the values of the χ^2 PDF are calculated. For example, for a 100 point f_γ and f_ρ, 10000 χ^2 PDF values are calculated without this convolution. With (5.17), however, only 200 PDF values are cal-

culated, though this depends on the required Δ for the composite distribution. This is a $50\times$ reduction in a complicated operation.

Consequently, (5.16) becomes:

$$f_{Y,SNR}[x] = step_{size\rho+\gamma} \sum_{\rho+\gamma=(\rho+\gamma)_{min}}^{(\rho+\gamma)_{max}} f_{\rho+\gamma}[\rho+\gamma] f_{Y_{s,\rho+\gamma}}[x]. \tag{5.18}$$

Equation (5.18) can then be used instead of (5.3) for all calculations when noise uncertainty is present.

5.4 Sorting Energy Detector Outputs

When attempting to predict the result of sorting the energy detector outputs, order statistics can be used. If N independent and identically distributed (i.i.d.) random variables, $(x_1, x_2, \cdots, x_N)$, each having the same PDF, $f(x)$, and CDF, $F(x)$, are sorted with respect to their magnitudes, order statistics allow the generation of the PDF of the n^{th} lowest valued variable, known as the n^{th} order statistic, $x_{n:N}$.

The PDF of the n^{th} order statistic of N variables, $f_{n:N}(x)$, is given by [110]:

$$f_{n:N}(x) = \frac{N!}{(N-n)!(n-1)!} F(x)^{n-1}(1-F(x))^{N-n} f(x). \tag{5.19}$$

In the noise-only case, let N_n denote the number of channels with no signal present. Thus, the PDF of the n^{th} noise-only variable, $x_{N,n:N_n}$, namely $f_{N,n:N_n}(x)$, is given by:

$$\begin{aligned} f_{N,n:N_n}(x) &= \frac{N_n!}{(N_n-n)!(n-1)!} \times \\ & F_N(x)^{n-1}(1-F_N(x))^{N_n-n} f_N(x), \end{aligned} \tag{5.20}$$

where $F_N(x)$ is the noise-only CDF at x and $f_N(x)$ is the noise-only PDF at x.

In the signal-and-noise case, we have N_{sig} channels with signals present. Thus, the PDF of the n^{th} signal-and-noise variable $x_{S,n:N_{sig}}$, denoted $f_{S,n:N_{sig}}(x)$, is given by:

$$\begin{aligned} f_{S,n:N_{sig}}(x) &= \frac{N_{sig}!}{(N_{sig}-n)!(n-1)!}(F_S(x))^{n-1} \times \\ &\quad (1-F_S(x))^{N_{sig}-n} f_{Y,SNR}(x), \end{aligned} \tag{5.21}$$

where $F_S(x)$ is the CDF of the energy detector output at x and $f_N(x)$ is the PDF of the energy detector output at x.

In Fig. 5.7 four signal-and-noise i.i.d. outputs, with $N_{coarse} = 1000$ and using the example SNR PDF from 4.4, are sorted by (5.21) and the result is shown along with the original PDF of the outputs.

For non-identically distributed variables, an alternative method for generating the sorted distributions is available [113]. However, non i.i.d. variables are not considered here as this method requires the computation of the permanent of an $N_a \times N_a$ matrix for each point in the distribution, where N_a is the number of energy detector outputs to be sorted and is equal to N_n and N_{sig} for the noise-only and signal-and-noise cases, respectively. The computation required for this is $O(2^{N_a} N_a)$ [114]. For example, if N_a = 9, then each point on the PDF would require 4.6k calculations, significantly increasing the time required for the simulation.

It will be shown in Section 5.9.1 that the error introduced into the final result due to the assumption of i.i.d. variables is small and the predicted optimum number of samples for the coarse detection N_{coarse} does not change significantly.

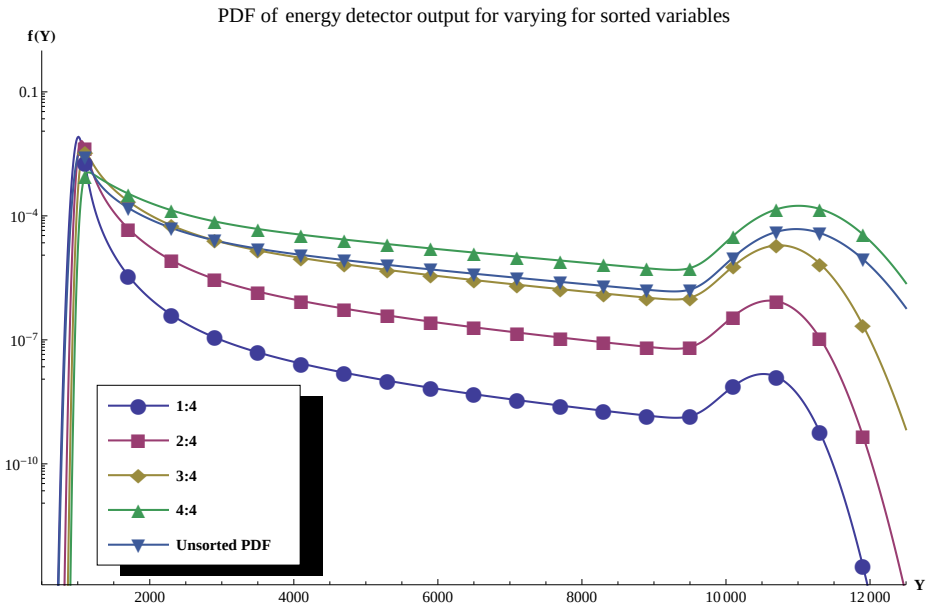

Figure 5.7: PDFs of Energy Detector Output When a Signal is Present and the Output is Sorted w.r.t. Magnitude

5.4.1 Comparing PDFs

Once the PDFs for the noise-only and signal-and-noise cases have been generated, they must be compared to find the sorting efficiency of the system.

The probability that a variable y_1, with PDF f_{y1}, will be smaller than a variable y_2, with PDF f_{y2}, assuming variables are independent, can be calculated as follows: The fraction of the signal PDF f_{y1} that is smaller than a value, k, is the probability that y_1 will be smaller than k.

Thus:

$$P_{y_1<k} = \int_{-\infty}^{k} f_{y1}(x)\mathrm{d}x. \tag{5.22}$$

For a sampled distribution where $f_{y1}[x]$ and $f_{y2}[x]$ = 0 for $x \leq 0$:

$$P_{y_1<k} = \Delta \sum_{j=1}^{k} f_{y1}[j \times stepsize], \tag{5.23}$$

where Δ is the step size of the summation.

Performing a weighted sum, based on the value of the second PDF, y_2, at these points, gives:

$$P_{y_1<y_2} = \Delta \sum_{k=1}^{N_{coarse}} P_{y_1<k} f_{y2}[k \times stepsize], \tag{5.24}$$

or

$$P_{y_1<y_2} = \Delta^2 \sum_{k=1}^{N_{coarse}} \sum_{j=1}^{k} f_{y1}[j] f_{y2}[k \times stepsize]. \tag{5.25}$$

In Fig. 5.8 the lowest of four noise-only outputs with the lowest of four signal-and-noise outputs, where the sample SNR distribution is used and $N_{coarse} = 1000$, are shown. Using (5.25) results in the probability of the lowest of four noise-only outputs being lower than the lowest of four signal-and-noise outputs being 98.1%. Comparing this to Monte-Carlo simulations run one million times, which produced a probability of 98.6% for the same event, shows that the method is accurate. If the Δ was reduced then this error would, on average, decrease further.

This method can be repeated to compare each of the signal-and-noise channels to the noise-only channels and generate a matrix, $\mathbb{P}$, where $\mathbb{P}$(i,j) is the probability that the i^{th} signal-and-noise output is greater than the j^{th} noise-only output, with all of the corresponding probabilities. This matrix will be used with Markov Chains to model the system with a view to calculating the system response.

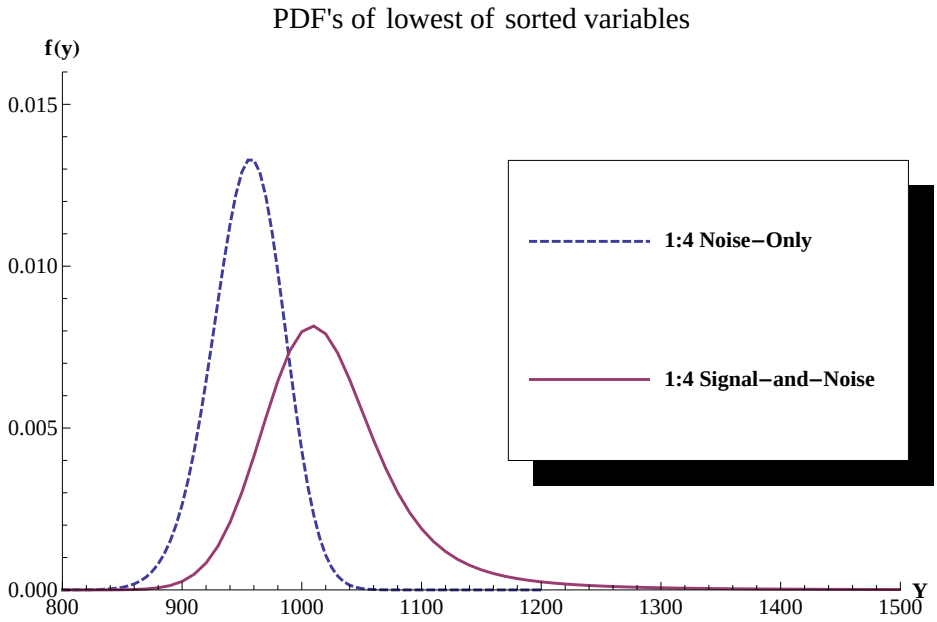

Figure 5.8: Comparison PDFs of Lowest of Four Sorted Energy Detector Outputs for Signal-And-Noise Case for Sample SNR Distribution, and Noise-Only Case for $N_{coarse} = 1000$

5.5 Time varying occupancies

Using the equations for timing in Section 3.7, the average probability of change in occupancy can be generated. As can be seen in Fig. 5.9 for a fine detector with approximately 200k samples, and a sampling rate of 1MHz, there is a low probability of interference ($\leq 10\%$) for $T_{ON} = T_{OFF} \geq 2$ seconds.

If this interference probability is sufficiently low that the Fine Detector's performance does not become significantly impaired then no changes are required to the fine detector. If this is not the case, then a weighting scheme, similar to that proposed in [43], could be used to increase performance.

Of greater impact on the CSFD is the fact that, as the number of detection attempts in-

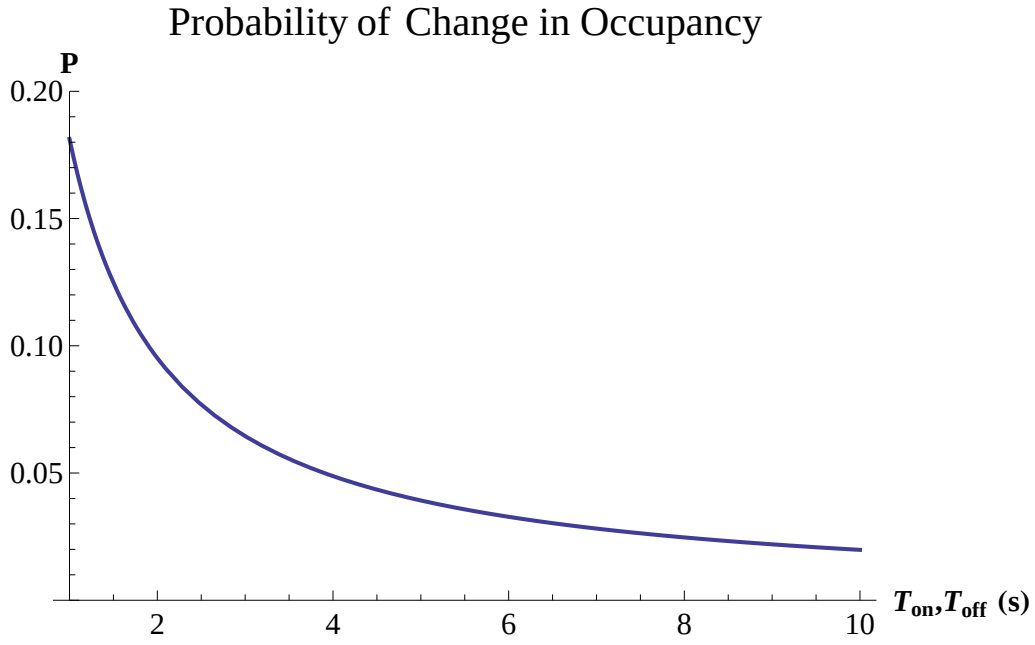

Figure 5.9: Probability of Change for t_{sense} =0.2s and Varying T_{ON}, T_{OFF}

creases, the accuracy of the sorting tends to decrease, in a channel with time varying occupancies. Here it is assumed that the probability of the channel switching occupancy twice during a sensing period is very low, i.e. that the $T_{ON} + T_{OFF} >> t_{sense}$. If the n^{th} detection attempt has a probability of scanning an occupied channel of P_{occ} (which is generated from (5.25)) then, with a time varying channel, the new probability is:

$$P_{occ,new} = P_{occ}(1 - P_{change,ON}[n-1]) + (1 - P_{occ})P_{change,OFF}[n-1], \qquad (5.26)$$

where $P_{change,OFF}[n-1]$ is the probability of change from occupied to unoccupied by the end of the $(n-1)^{th}$ detection attempt and uses T_{ON} as the mean channel time and $P_{change,OFF}[n-1]$ is the probability of change from unoccupied to occupied by the end of the $(n-1)^{th}$ detection attempt and uses T_{OFF} as the mean off time. Note that only the

probability of the detector changing before the n^{th} detection attempt is used here, as this will lead to the decision becoming inaccurate.

In Fig. 5.10 the new occupancy probabilities are shown for the n^{th} detection attempt, $2 \leq n \leq 6$, for varying T_{ON} and T_{OFF}. The sensing time, t_{sense}, is now $0.2 \times (n-1)$.

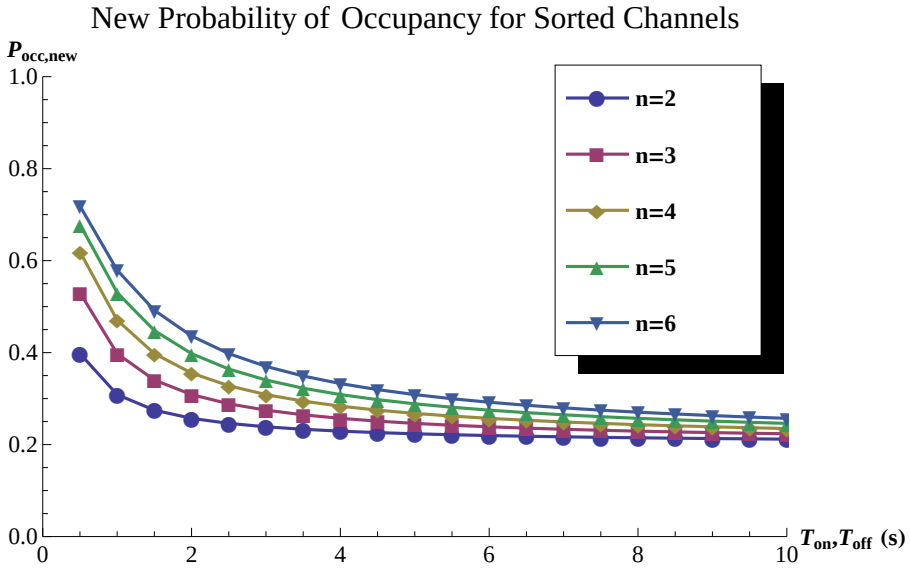

Figure 5.10: $P_{occ,new}$ for t_{sense} =0.2s, P_{occ} =0.8 and Varying T_{ON}, T_{OFF}

For $T_{ON} = T_{OFF} \geq 10$s, there is little change in the occupancy probability, i.e. for n=6 and $T_{ON}, T_{Off} = 10$s, $P_{occ,new} = 0.74$. However, the T_{ON} and T_{OFF} times will not be equal when the average occupancy, θ, is not equal to 50%.

The average occupancy, θ, is a function of the average "ON" and "OFF" times and can be written as

$$\theta = \frac{T_{ON}}{T_{ON} + T_{OFF}}. \tag{5.27}$$

If $T_{ON} + T_{OFF}$ is set to a specific value, then it is easy to find the appropriate values for T_{ON} and T_{OFF} for a specific θ. In Fig. 5.11 this is shown.

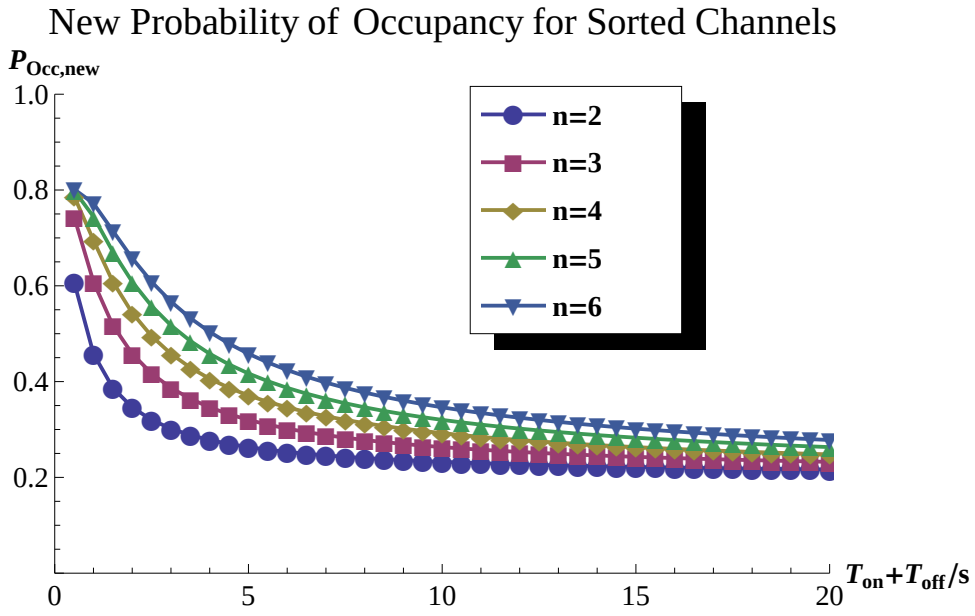

Figure 5.11: $P_{occ,new}$ for t_{sense} =0.2s, P_{occ} =0.2 and Varying $T_{ON} + T_{OFF}$ and θ= 60%

For the third detection attempt at an average occupancy of 60% and $T_{ON}+T_{OFF}$ of 5s, the new occupancy probability equals 0.25 (the expected value is 0.2 without a time varying occupancy).

In this system it is unlikely that the detector will require six detections, indeed it will be shown later that the average number of detection attempts required by coarse fine sensing is approximately two. Thus, small variations in the probability of occupancy will not change the model significantly. If the occupancy is time varying with a sufficiently small $T_{ON} + T_{OFF}$, then the model can account for this by using (5.26).

5.6 Markov Chain Model

A Markov Chain can be used to model the system once the probability matrix, $\mathbb{P}$, has been generated [98]. In real systems the secondary detector will have a finite P_{fa} and P_{md}. Thus, to have an accurate model of the system, it must include these probabilities of failure.

The sorting performed by the coarse detector does not depend on the secondary detector. Therefore, the probability matrix $\mathbb{P}$ is independent of the secondary detector used.

The Markov model is somewhat complicated, as the system cannot assume that the fine detector will detect the spectrum opportunity on the first attempt. The system must compare all the noise-only energy detector outputs to the signal-and-noise energy detector outputs.

The Markov model has a significant number of possible paths. In Fig. 5.12 the transition probabilities are shown for a single transient state, $S_{m,n}$.

State $S_{m,n}$ is the decision between the m^{th} signal-and-noise energy detector output and n^{th} noise-only energy detector output. From state $S_{m,n}$ there are four possible paths.

The system will transition from $S_{m,n}$ to $S_{m+1,n}$ with a transition probability of $P_{T1}(m,n)$. It will transition from $S_{m,n}$ to $S_{m,n+1}$ with a transition probability of $P_{T2}(m,n)$.

The system will transition from $S_{m,n}$ to the missed detection absorbing state with a transition probability of $P_{Int}(m,n)$. It will transition from $S_{m,n}$ to the free channel detected absorbing state with a transition probability of $P_{FC}(m,n)$. Equations (5.30)-(5.33) below specify how these probabilities are calculated.

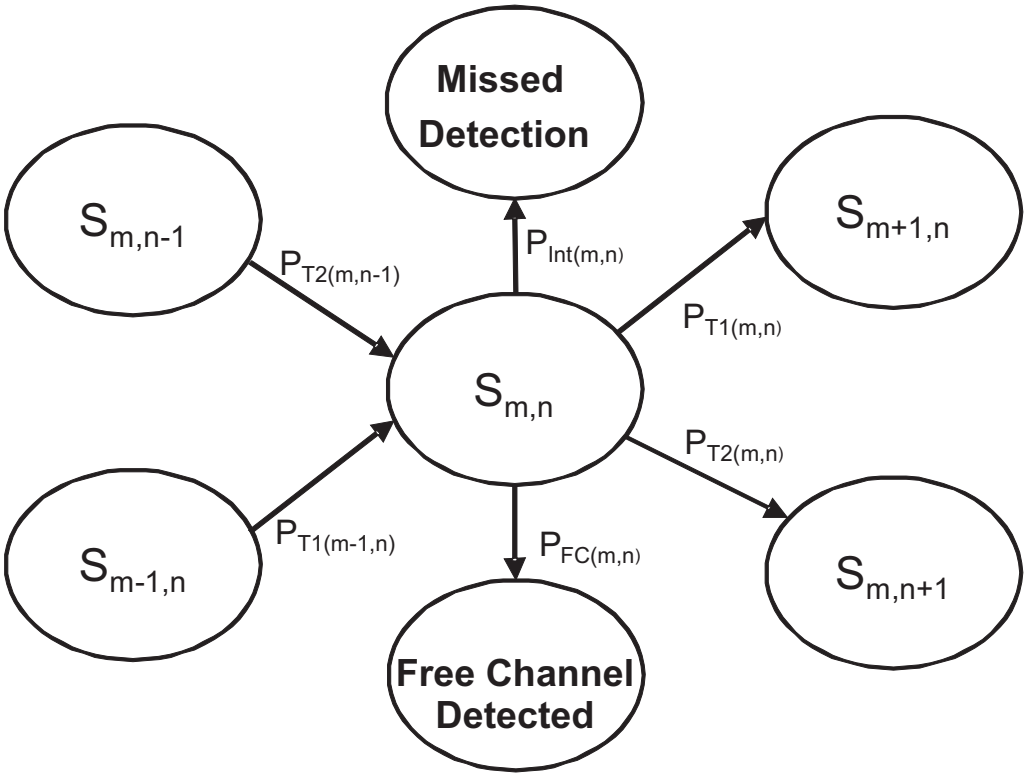

Figure 5.12: Markov Chain for CSFD and Naive Models

5.6.1 Naive Detector

For a Naive detector with a non-ideal secondary detector, the probability that the channel to be scanned in state $S_{m,n}$ is unoccupied, $P_N(m,n)$, is given by:

$$P_N(m,n) = \begin{cases} 0 & \text{if } n > N_{Noise} \\ 1 & \text{if } m > N_{Sig} \\ \frac{N_{noiser}}{N_{chr}} & \text{otherwise} \end{cases}, \tag{5.28}$$

where N_{noiser} is the number of remaining noise-only channels and N_{chr} is the total number of channels remaining.

There is no possibility of both $n > N_{Noise}$ and $m > N_{Sig}$ being simultaneously true as

this would require more detection attempts than channels. Since state $S_{m,n}$ corresponds to having $m-1$ signal-and-noise channels and $n-1$ noise-only channels scanned previously, (5.28) becomes:

$$P_N(m,n) = \begin{cases} 0 & \text{if } n > N_{Noise} \\ 1 & \text{if } m > N_{Sig} \\ \frac{N_{noise}-n+1}{N_{ch}-n-m+2} & \text{otherwise} \end{cases}. \tag{5.29}$$

The transition probabilities for the Markov chain are defined, where $\bar{P}_{md}$ is the average P_{md} over the SNR distribution, as:

- P_{T1} is the probability of correctly deciding that a signal is present, requiring the system to continue scanning from state $S_{m+1,n}$, given by:

$$P_{T1}(m,n) = (1 - P_N(m,n))(1 - \bar{P}_{md}). \tag{5.30}$$

- P_{T2} is the probability of incorrectly deciding that a signal is present, requiring the system to continue scanning from state $S_{m,n+1}$, given by:

$$P_{T2}(m,n) = P_N(m,n)P_{fa}. \tag{5.31}$$

- P_{Int} is the probability of incorrectly deciding that no signal is present, thus causing harmful interference to the primary user, given by:

$$P_{Int}(m,n) = (1 - P_N(m,n))\bar{P}_{md}. \tag{5.32}$$

- P_{FC} is the probability of correctly deciding that no signal is present, thus finding a

spectrum opportunity, given by:

$$P_{FC}(m,n) = P_N(m,n)(1 - P_{fa}). \tag{5.33}$$

P_{fa} will remain constant at the chosen value but, for the fine detector architectures considered here, P_{md} will depend on the PDF of the SNR.

Equations (5.29)-(5.33) allow the Markov Chain matrices to be populated for the Naive system, where there is a non-ideal secondary detector. The transitional matrix $\mathbb{Q}$ [98] is a square matrix of order $\frac{n(n+1)}{2}$. For $N_{ch} = 10$, $\mathbb{Q}$ is a 55×55 matrix. For an absorbing Markov chain the fundamental matrix, $\mathbb{N}$, is defined as [98]:

$$\mathbb{N} = (\mathbb{I} - \mathbb{Q})^{-1}, \tag{5.34}$$

where $\mathbb{I}$ is an identity matrix of corresponding size.

The average number of steps required to reach an absorbing state, starting from the i^{th} state, is then given by $\bar{\mathbb{N}}_{steps}[i]$:

$$\bar{\mathbb{N}}_{steps}[i] = \mathbb{N}\xi[i], \tag{5.35}$$

where ξ denotes the row sum of a matrix.

Since each step corresponds to a detection attempt, the average total number of samples required for the detector, $\bar{N}_{T,Naive}$, is given by:

$$\bar{N}_{T,Naive} = \bar{\mathbb{N}}_{steps}[i] N_{fine}, \tag{5.36}$$

where N_{fine} is the number of samples required for fine detection.

For an energy detector to guarantee a P_{fa} and P_{md} of 10% at an SNR of -21dB, it requires

Table 5.1: The increase in number of detection attempts required for increasing occupancy for a Naive Detector

θ	$N_{T,Naive}$ PDF Method	$N_{T,Naive}$ Monte-Carlo Method
0.1	254k	255k
0.2	282k	283k
0.3	317k	319k
0.4	363k	364k
0.5	423k	425k
0.6	507k	509k
0.7	633k	636k
0.8	841k	844k
0.9	1230k	1236k

$N_{fine} = 209k$ samples. $N_{T,Naive}$ was generated for these values and compared to Monte Carlo simulations to verify accuracy. The results are shown in Table 5.1.

The results for the Markov Chain analysis match very closely with those found using Monte-Carlo simulations. Even at 90% occupancy, where θ denotes the occupancy and is equal to 0.9 in this case, the relative error is still less than 0.6%.

5.6.2 Coarse-Sorting Fine Detector

For the CSFD detector, using a non-ideal secondary detector, each transient state in the chain corresponds to the comparison of a signal-and-noise energy detector output PDF to a noise-only energy detector output PDF. Since the energy detector outputs are sorted, the P_{md} for the lowest signal-and-noise will generally not be equal to the average P_{md} for the SNR distribution. Consequently, the effects of the ordering must be accounted for.

The SNR of the outputs can be thought to be ordered on average, though not in all cases. It is possible that the n^{th} signal-and-noise energy detector output could have an SNR that was greater than that of the $(n+1)^{th}$ signal-and-noise energy detector output. On

average, however, the n^{th} signal-and-noise energy detector output will have a lower SNR than successive outputs.

For four signal-and-noise channels with the sample SNR distribution, the signals with the lowest energy detector output have a P_{md} of 0.206%. Using the sorting method, this is predicted to be 0.232%. Since the P_{md} is so low this accuracy is deemed sufficient for detector speed calculations. It will, however, influence the overall P_{md} calculations and reduce the accuracy of the calculations.

It is, therefore, deemed acceptable to use order statistics upon the SNR distribution to get an estimate of the distribution of the n^{th} signal-and-noise energy detector output, for the purposes of calculating P_{md}.

The sorted SNR PDF, $f_{SNR,n:N_{sig}}[x]$, is given by:

$$\begin{aligned} f_{SNR,n:N_{sig}}[x] &= \frac{N_{sig}!}{(N_{sig}-n)!(n-1)!} F_{SNR}[x]^{n-1} \times \\ &\quad (1-F_{SNR}[x])^{N_{sig}-n} f_{SNR}[x]. \end{aligned} \tag{5.37}$$

Thus, the average P_{md} associated with the n^{th} detection attempt, $\bar{P}_{md}(n)$, is:

$$\bar{P}_{md}(n) = \Delta \sum_{\gamma=SNRmin}^{SNRmax} f_{SNR,n:N_{sig}}[\gamma] P_{md}(\gamma), \tag{5.38}$$

where $P_{md}(\gamma)$ is the probability of missed detection at an SNR of γ and Δ is the step size of the summation or the "resolution" of the SNR PDF.

When the comparisons of the PDFs are completed using (5.25), the result is the probability that the n^{th} noise-only energy detector output PDF will be smaller than the m^{th} signal-and-noise energy detector output PDF, denoted $P_{orig}(m,n)$.

For state $S_{m,n}$, $P_a(m,n)$ is the probability that the n^{th} noise-only coarse detector output will be smaller than the m^{th} signal-and-noise coarse detector output, **given that** the detector has reached state $S_{m,n}$.

Similar to (5.29), $P_a(m,n)$ can be expressed as :

$$P_a(m,n) = \begin{cases} 0 & \text{if } n > N_{Noise} \\ 1 & \text{if } m > N_{Sig} \\ P_{orig}(m,n) \mid S_{m,n} & \text{otherwise} \end{cases} . \tag{5.39}$$

For state $S_{m,n}$, i.e. the comparison of the m^{th} signal-and-noise Coarse Detector output PDF to the n^{th} noise-only energy detector output PDF, the transitional probabilities generated are:

$$P_{T1}(m,n) = (1 - P_a(m,n))(1 - P_{md}(m)), \tag{5.40}$$

$$P_{T2}(m,n) = P_a(m,n)P_{fa}, \tag{5.41}$$

$$P_{Int}(m,n) = (1 - P_a(m,n))P_{md}(m), \tag{5.42}$$

$$P_{FC}(m,n) = P_a(m,n)(1 - P_{fa}). \tag{5.43}$$

From Fig. 5.12 there are two possible paths to state $S_{m,n}$. The probability of reaching state $S_{m,n}$ is the sum of these two probabilities, i.e.:

$$P(m,n) = P_{p1}(m,n) + P_{p2}(m,n). \tag{5.44}$$

The two associated probabilities must be treated separately and the results weighted by the probabilities of each path respectively (P_{p1} and P_{p2}) to give the correct result.

5.6.2.1 Path 1

In state $S_{m-1,n}$, the path from state $S_{m-1,n}$ to state $S_{m,n}$, which occurs when the fine detector correctly declares a signal present, has a transitional probability of $P_{T1}(m-1,n)$. The probability of this path being chosen, $P_{p1}(m,n)$, is the transitional probability of the previous state to the current state, $P_{T1}(m-1,n)$, multiplied by the probability of reaching state $S_{m-1,n}$, denoted $P(m-1,n)$, i.e.

$$P_{p1}(m,n) = P(m-1,n)P_{T1}(m-1,n). \tag{5.45}$$

In the previous detection attempt the $(m-1)^{th}$ signal-and-noise channel was selected for fine scanning with probability $1 - P_{orig}(m-1,n)$ (event A). The channel was then correctly declared occupied with probability $1 - P_{md}(m)$. The probability that the m^{th} signal-and-noise coarse detector output is smaller than the n^{th} noise-only energy detector output is $1 - P_{orig}(m,n)$ (event B).

Since event A has already occurred then the probability of event B occurring is $1 - P_{a,p1}(m,n)$ where,

$$P_{a,p1}(m,n) = 1 - P(B \mid A) = 1 - \frac{P(B \cap A)}{P(A)}. \tag{5.46}$$

If the m^{th} signal-and-noise energy detector output is smaller than the n^{th} noise-only energy detector output, then the $(m-1)^{th}$ signal-and-noise coarse detector output must also be smaller, as it is smaller than the m^{th} signal (due to the ordering). Therefore, $P(B \cap A) = P(B)$ and $P_{a,p1}(m,n)$ becomes:

$$P_{a,p1}(m,n) = 1 - \frac{P(B)}{P(A)} = 1 - \frac{1 - P_{orig}(m,n)}{1 - P_{orig}(m-1,n)}. \tag{5.47}$$

5.6.2.2 Path 2

In state $S_{m,n-1}$, the path from state $S_{m,n-1}$ to state $S_{m,n}$, which occurs where the fine detector incorrectly declares a signal present, has a transitional probability of $P_{T2}(m, n-1)$. Therefore, similar to (5.45), the probability of path 2 being chosen is given by:

$$P_{p2}(m,n) = P(m,n-1)P_{T2}(m,n-1). \tag{5.48}$$

In the previous detection attempt the $(n-1)^{th}$ noise-only channel was selected for fine scanning with probability $P_{orig}(m,n-1)$ (event C). The channel was then incorrectly declared occupied with probability P_{fa}.

The probability that the m^{th} signal-and-noise coarse detector output is lower than the n^{th} noise-only energy detector output remains equal to $1 - P_{orig}(m,n)$ (event D). Since event C has already occurred, the probability of event D occurring is $1 - P_{a,p2}(m,n)$ where,

$$P_{a,p2}(m,n) = 1 - P(D \mid C) = 1 - \frac{P(D \cap C)}{P(C)}. \tag{5.49}$$

Similar to (5.47), it can be shown that,

$$P_{a,p2}(m,n) = \frac{P_{orig}(m,n)}{P_{orig}(m,n-1)}. \tag{5.50}$$

Thus, P_a becomes:

$$P_a(m,n) = \frac{P_{a,p1}(m,n)P_{p1}(m,n) + P_{a,p2}(m,n)P_{p2}(m,n)}{P(m,n)}. \tag{5.51}$$

Using (5.47) and (5.50), $P_a(m,n)$ is given by

$$\begin{aligned} P_a(m,n) &= (1-\frac{1-P_{orig}(m,n)}{1-P_{orig}(m-1,n)})\frac{P_{p1}(m,n)}{P(m,n)} + \\ &\quad (\frac{P_{orig}(m,n)}{P_{orig}(m,n-1)})\frac{P_{p2}(m,n)}{P(m,n)}. \end{aligned} \tag{5.52}$$

It must be noted that the values for $S_{m,n}$ depend on both previous states, $S_{m-1,n}$ and $S_{m,n-1}$. The values must be generated consecutively, starting in the initial state $S_{1,1}$, where the first detection attempt is considered, moving on to $S_{1,2}$ and $S_{2,1}$ and continuing until all the transition probabilities for all possible states have been generated. For the initial state $S_{1,1}$ we have: $P(1,1) = 1$, i.e. the detector has to start in this state, and $P_a(1,1) = P_{orig}(1,1)$ (since there are no previous decisions to influence this).

It is then possible to populate the matrices for the CSFD Markov Chain and to use (5.34) and (5.35) to generate the average number of detection attempts required by the CSFD detector, $\bar{\mathbb{N}}_{T,CSFD}(N_{coarse},\theta)$, for a number of coarse detector samples, N_{coarse}, and occupancy, θ:-

$$\bar{\mathbb{N}}_{T,CSFD}(N_{coarse},\theta) = \bar{\mathbb{N}}_{steps}(N_{coarse},\theta)N_{fine} + N_{coarse}N_{ch}. \tag{5.53}$$

For an energy detector to guarantee a P_{fa} and P_{md} of 10% at an SNR of -21dB, N_{fine} is 209k samples. $\bar{N}_{T,CSFD}(N_{coarse},\theta)$ was generated for these values and compared to Monte Carlo simulations to verify accuracy. N_{coarse} was set at 1k samples. The results are shown in Table 5.2. The results of this comparison for the Markov chain match very closely indeed to the Monte-Carlo simulations, (for $\theta = 0.9$ the relative error is still less than 0.5%).

Table 5.2: The increase in number of detection attempts required for increasing occupancy for a CSFD Detector

θ	$N_{T,CSFD}$ PDF Method	$N_{T,CSFD}$ Monte-Carlo Method
0.1	243k	244k
0.2	247k	248k
0.3	253k	254k
0.4	261k	262k
0.5	272k	273k
0.6	289k	290k
0.7	320k	320k
0.8	387k	388k
0.9	622k	624k

5.7 Expected Speed Gain and Interference Rates

The average number of samples required for detection, under the conditions specified here, can be used to find the optimum number of coarse detector samples and, thus, the maximum speed gain possible for a CSFD detector over a Naive detector. For an occupancy θ and N_{coarse} coarse detector samples, equation (5.53) gives the expected number of samples required for detection for a CSFD detector.

If θ is known then a search can be performed to find the optimum number of samples. This can be done graphically by generating the full curve for a range of N_{coarse} values and finding the value of N_{coarse} that gives the minimum $\bar{N}_{T,CSFD}(N_{coarse}, \theta)$.

Another option is to search numerically, evaluating $\bar{N}_{T,CSFD}(N_{coarse}, \theta)$ only for specific values and attempting to find the minimum. For the SNR PDFs considered here, $\bar{N}_{T,CSFD}$, is smooth and unimodal w.r.t N_{coarse}, as can be seen in Fig. 5.14. Thus, it is expected that a simple search algorithm can be used to find the minimum.

For an unknown θ, the average $\bar{N}_{T,CSFD}(N_{coarse}, \theta)$ over the occupancy must be found. If the distribution is not known then it must be assumed to be uniform. If this is not the

case, then the individual values of $\bar{N}_{T,CSFD}(N_{coarse}, \theta)$ must be scaled by the PDF of θ, denoted f_θ.

The average $\bar{N}_{T,CSFD}$ for a coarse fine detector, $\bar{N}_{avg,CSFD}$, can be expressed as:

$$\bar{N}_{avg,CSFD}(N_{coarse}) = \int_0^1 f_\theta(\theta)\bar{N}_{T,CSFD}(N_{coarse}, \theta)\mathrm{d}\theta. \tag{5.54}$$

The average number of samples used for a Naive detector, $\bar{N}_{T,Naive}$, can be expressed as:

$$\bar{N}_{avg,Naive} = \int_0^1 f_\theta(\theta)\bar{N}_{T,Naive}(\theta)\mathrm{d}\theta. \tag{5.55}$$

Consequently, the speed gain of the system is expressed as:

$$\bar{G}_s(N_{coarse}) = \frac{\bar{N}_{avg,Naive}}{\bar{N}_{avg,CSFD}(N_{coarse})}. \tag{5.56}$$

The Markov Chain can also be used to generate a matrix of the variance in the number of steps required to reach a decision, $\mathbb{V}_{steps}$. It is given by [98]:

$$\mathbb{V}_{steps} = (2\mathbb{N} - \mathbb{I})\bar{\mathbb{N}}_{steps} - \bar{\mathbb{N}}_{steps,sq}, \tag{5.57}$$

where $\bar{\mathbb{N}}_{steps,sq}$ is the square of each of the entries in the $\bar{\mathbb{N}}_{steps}$ matrix or

$$\bar{\mathbb{N}}_{steps,sq}[i,j] = \bar{\mathbb{N}}_{steps}[i,j]^2. \tag{5.58}$$

$Var[N_{steps}]$ is then a vector of length equal to the number of states and $Var[N_{steps}][i]$ gives the variance of the number of steps required to reach a decision, starting in the i^{th} state. This then allows the variance of the number of samples required for CSFD to be found.

It is also possible to find the probability of incorrect operation. The probability matrix, $\mathbb{B}$, of entering an absorbing state is given by:

$$\mathbb{B} = \mathbb{N}\mathbb{R}, \tag{5.59}$$

where $\mathbb{R}(i, j)$ is the matrix term that governs the probability of transition from the i^{th} transient state to the j^{th} absorbing state. The first entry in the $\mathbb{B}$ matrix corresponds to the probability of entering the first absorbing state.

The probability of causing interference is given by the probability of deciding an occupied channel is unoccupied and using it for transmission. In the Markov chain studied here, this occurs if the system enters the missed detection absorbing state. Similarly, there is a probability of missing a free channel. If there is a free channel available, and the radio declares it occupied, then there is a missed opportunity. This occurs if the system enters the missed opportunity state. The appropriate probability can be found by selecting the relevant entry in $\mathbb{B}$.

5.8 Fading

If the system is subjected to fading then the SNR PDF is modified by the appropriate fading distribution. Here, attention is confined to slow fading. Slow fading occurs when the channel impulse response changes slowly, such that the SNR is constant within each detection attempt but may vary between detection attempts. Not that shadow fading is not considered here.

For example, with Rayleigh fading [55], the PDF of γ is:

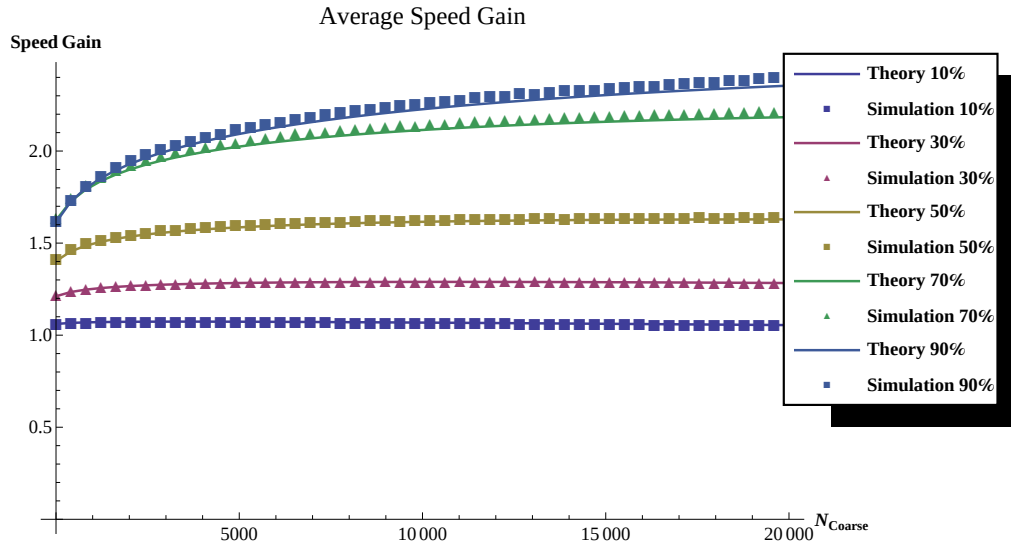

Figure 5.13: Verification of Model Under Fading for Non-Ideal Secondary Detector for different levels of occupancy

$$f_{fad}(\gamma, \bar{\gamma}) = \frac{1}{\bar{\gamma}} e^{-\gamma/\bar{\gamma}}; \gamma \geq 0, \tag{5.60}$$

where $\bar{\gamma}$ is the average SNR. Distributions for Ricean and other channels, as well as methods for finding the probabilities of missed detection for a signal under fading, can be found in [55].

Replacing the single value of $\bar{\gamma}$ in (5.60) with the user defined SNR distribution, $f_{SNR}(\bar{\gamma})$, and averaging, the new SNR distribution under fading can be calculated by

$$f_{SNRfaded}(\gamma) = \sum_{\bar{\gamma}} f_{SNR}(\bar{\gamma}) f_{fad}(\gamma, \bar{\gamma}). \tag{5.61}$$

This new SNR distribution can then be used with the PDF method to generate the design curves without any further changes to the model.

The fine detector will need to be modified to account for the fading also. An energy detec-

tor subjected to Rayleigh fading with an average SNR of -21dB requires approximately 4.8M samples to detect a signal with a P_{md} of 10% and a P_{fa} of 10%. This value can be found by averaging the P_{md} over the SNR distribution of a Rayleigh faded signal at an average SNR of -21dB [55]. If an energy detector is used, then these probabilities can be generated with $f_{SNRfaded}(\gamma)$ replacing $f_{SNR}(\gamma)$ in equation (5.37). Verification of the model under Rayleigh fading for a non-ideal secondary detector is shown in Fig. 5.13.

5.9 Results

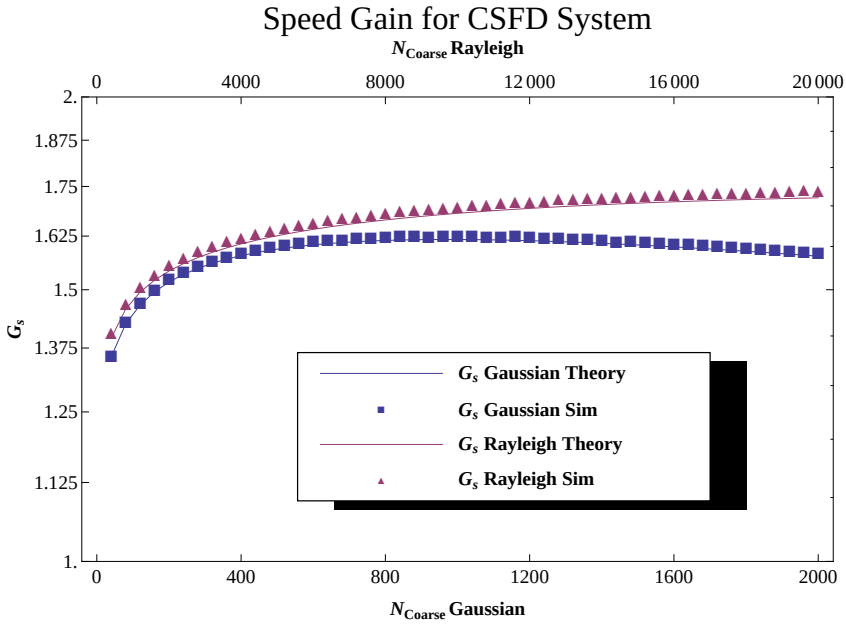

Figure 5.14: Speed Gain G_s vs Number of Samples for Varying N_{coarse} for Rayleigh Fading and Gaussian (No Fading) Noise.

Shown in Fig 5.14 is the predicted speed gain performance under Rayleigh fading. Also shown is the performance without fading, i.e. for a Gaussian channel. Note that, although the gain possible for both conditions is roughly equal, the number of samples for detection

under Rayleigh fading is significantly larger.

Fig. 5.15 shows the P_{fa} and P_{md} for Gaussian and Rayleigh channels. The P_{md} matches well, though the sorting effects are not properly calibrated (as discussed in section 5.6.2) the magnitude of the error is small. The P_{fa} does not match as well, though, as the average P_{fa} is approximately 0.01, the effects of this error are small.

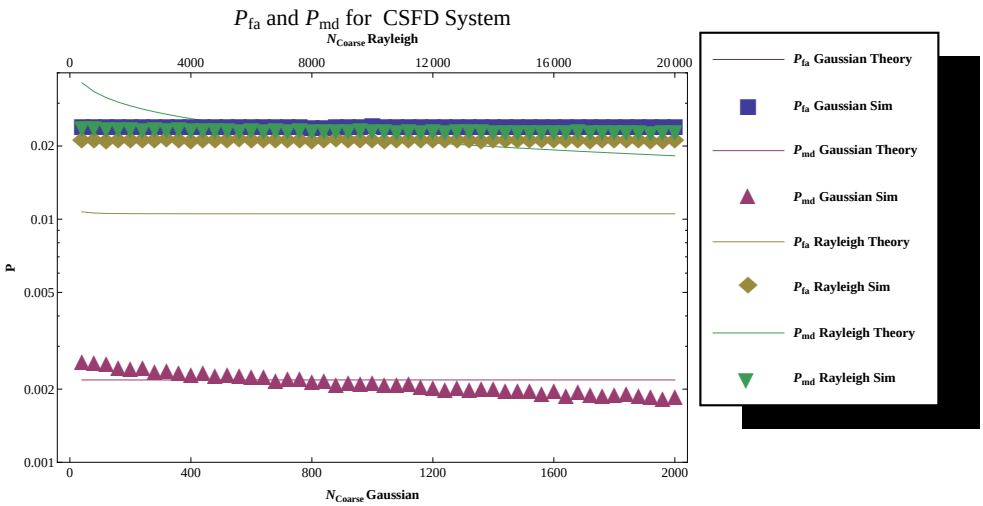

Figure 5.15: P_{fa} and P_{md} Under Gaussian and Rayleigh Channels

The PDF Method was significantly quicker than the Monte-Carlo Method. The PDF method took approximately four hours to generate the design curves above [i]. The Monte-Carlo Method required approximately two weeks on the same machine. Therefore, the PDF is approximately 84 times faster.

5.9.1 Non-I.I.D. SNR Distributions

As stated in section 5.4 the assumption of i.i.d. for SNRs of the signals to be sorted is made to reduce the complexity of the sorting calculations. Now that the full system model

[i] Single machine single processor at 2.53GHz and 3GB of RAM

is available, the expected speed gain for non-i.i.d. channels is investigated.

Fig 5.16 shows a sample SNR PDF for ten channels together, where the channels have different SNR distributions. The channels have SNR distributions that are Gaussian in shape with different means and standard deviations. These are summarised in Table 5.3.

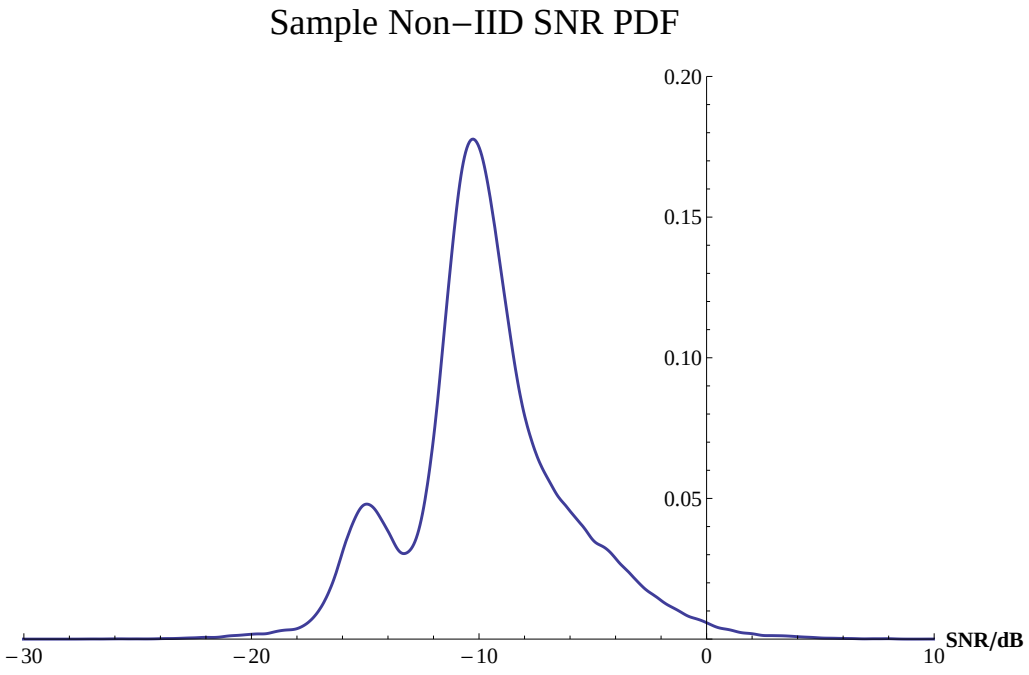

Figure 5.16: Overall SNR PDF for all Ten Channels

The sample distributions were used in three ways to test the i.i.d. assumption. First, the overall distribution of the SNRs were used with the model developed here to predict performance assuming that all the SNR distributions were i.i.d.

The second test was a Monte-Carlo simulation where the SNR was generated for each of the bands and then used individually on the channels. This result shows the average speed gain for the system where the SNR distributions were independent but not identically distributed.

Table 5.3: The example SNR distributions used for test

Channel	Mean \dB	Standard Deviation \dB
1	-9.9	5
2	-11	0.8
3	-8	5
4	-15	1
5	-7	2.6
6	-10	1
7	-6	3
8	-10	1
9	-9.5	1.5
10	-11.1	2

The final test case generated the SNRs for each distribution but the channels were selected randomly from the list, with the possibility that two channels could have the same SNR. This result shows the average speed gain for the system where the SNR distributions for some of the channels were not independent and the channels were not identically distributed.

The speed gains for all three cases are shown in Fig. 5.17. The maximum speed gain is correctly predicted by the PDF method using the i.i.d. assumption to within 0.03% of both of the other cases ($\approx 1.74\times$). The maximum is located for the same number of coarse samples for the i.i.d. assumption and for the case where the channels are independent but with non-identical distributions ($N_{coarse} = 2.1k$ samples) but 100 samples higher for the case where the channels can be dependent also ($N_{coarse} = 2.2k$ samples). This result shows that the i.i.d. assumption is valid and that the improvement of CSFD's performance over a Naive detector does not change dramatically.

There is a possibility for further improvement upon the CSFD. After the first few detection attempts it is very unlikely that there are any free channels remaining. If the CSFD

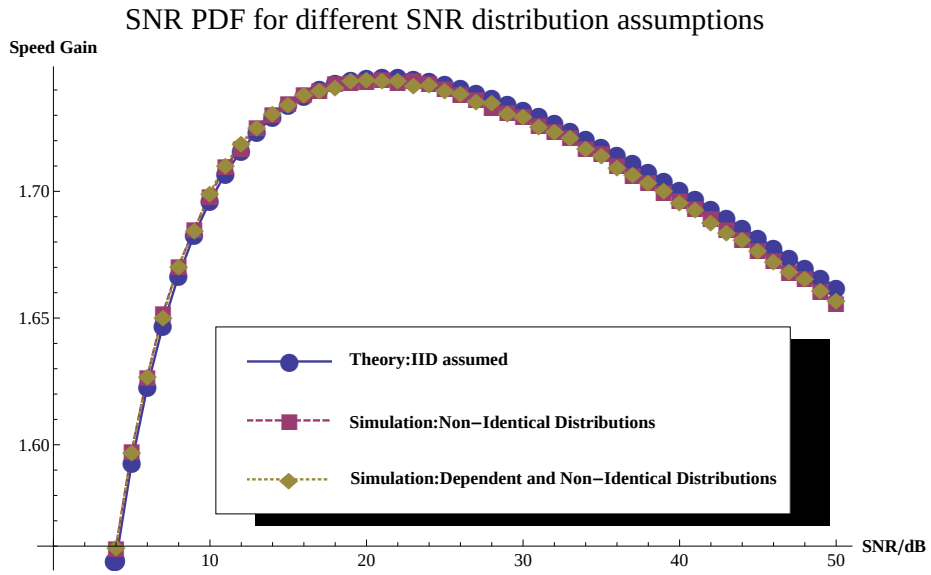

Figure 5.17: Speed Gain G_s vs Number of Samples for Varying N_{coarse} for Different SNR Distribution Assumptions

truncates its search then, at high occupancies, the performance can be improved. This is investigated further in Chapter 7, where a hybrid detector uses a similar principle to increase performance.

5.10 Conclusions

This work has presented a method enabling the modelling of performance for coarse fine spectrum sensing. By the use of Markov chains and order statistics, accurate estimation of system performance is possible. This method is significantly faster (approximately 100 times) than Monte-Carlo simulations.

The mismatch in the P_{fa} does not change the overall speed gain significantly. In addition,

the system has its performance at -21dB guaranteed by the fine detector alone.

The use of a SNR distribution is a more accurate method of estimating system performance in a real environment. The method shown here allows the easy inclusion of user defined SNR distributions. Provided the PDF of the distribution is available then this approach can be used to analyse it. The Monte-Carlo method would require the ability to generate a variable with this distribution.

The channels have been approximated as i.i.d. here to reduce computational complexity. It was shown that this approximation does not significantly impact on the accuracy of the model presented here, with less than 0.03% of a difference in the predicted maximum speed gain.

This work also assumes a slowly varying channel, where the occupancies are not varying during the detection attempt. If the occupancies are varying with time, then the Markov model will have to be modified to account for this. As DTV channels are considered in this work it is deemed reasonable make the assumption of fixed occupancies during detection.

Using this work design curves can be generated quickly, allowing the system to be optimised. These curves allow the designer to pick the value for N_{coarse} that gives the best performance under the expected signal conditions.

The value for $\bar{N}_{avg,CSFD}(N_{coarse})$ does not change significantly for small changes in N_{coarse}. Therefore, it might be possible to select a value for N_{coarse} that is appropriate for a wide range of distributions. This is investigated further in Chapter 7.

The truncation of a search when the system has a high degree of confidence that a channel is occupied is considered in the next chapter. The CDFD is a detector that, instead of sorting the channels, makes a decision on whether the channel is occupied or whether a more accurate scan should be performed. This also allows multiple coarse stages to be

used, as will now be discussed.

6

Coarse Deciding Fine Detector

6.1 Introduction

6.1.1 Multi-Resolution Architecture

As stated in the previous chapter, CF spectrum sensing uses a quick scan to generate information about a band and the energy levels in that band that can be used to optimise the fine scan. A coarse scan can provide information about the spectrum that can reduce

the usage of the fine detector and, thus, the total detection time. The previous chapter presented a method enabling the optimum number of samples to be found for a single coarse scan, for a user-specified SNR distribution. However, if a more accurate coarse scan was used between the coarse and fine ones, could it reduce the total detection time even further? Indeed, is there an optimum number of scans to be used?

Obviously, repeatedly sorting the channels would result in no advantage to using multiple detection attempts, or multiple regions. The detector must be able to remove some of the candidate channels from consideration at each detection attempt for multiple regions to be of advantage. To be able to reliably detect occupied channels at a low SNR means that reliably deciding a channel is unoccupied, when using a small number of samples, is difficult. However, reliably deciding a channel is occupied with a small number of samples is possible if the SNR is sufficiently high.

The detector operates as follows: firstly, after each detection attempt it decides if there is a signal present or not. If the detector decides that there is a signal present then it can declare the channel occupied. If the detector decides that there is no detectable signal present, then the next, finer, detection attempt is used. A channel is only declared unoccupied after the final detection attempt has failed to detect a signal. This final detection attempt can then decide that a signal at -21dB is not present, with the required probabilities of false alarm and missed detection. This again decouples the safety of the system from its speed performance.

A sample flow chart for the Multi-Resolution sensor is shown below in Fig. 6.1. The system has N_R sensing regions, or N_R different sensing attempts. For the i^{th} region, or sensing attempt, the system takes N_i samples and the energy detector is run on these samples for an output Y_i. Samples are re-used such that the i^{th} threshold is set for the total number of samples taken up to that point, not just N_i, and the sum of all Y_j $(1 \leq j \leq i)$

outputs is compared to the threshold to determine if the signal is present. If the threshold is not exceeded, then the system moves through succesive regions until all of the detection attempts have been completed. The channel is then declared unoccupied if the threshold has not been exceeded. If, however, the threshold is exceeded at any point then the system declares the channel occupied and moves to the next channel to be scanned.

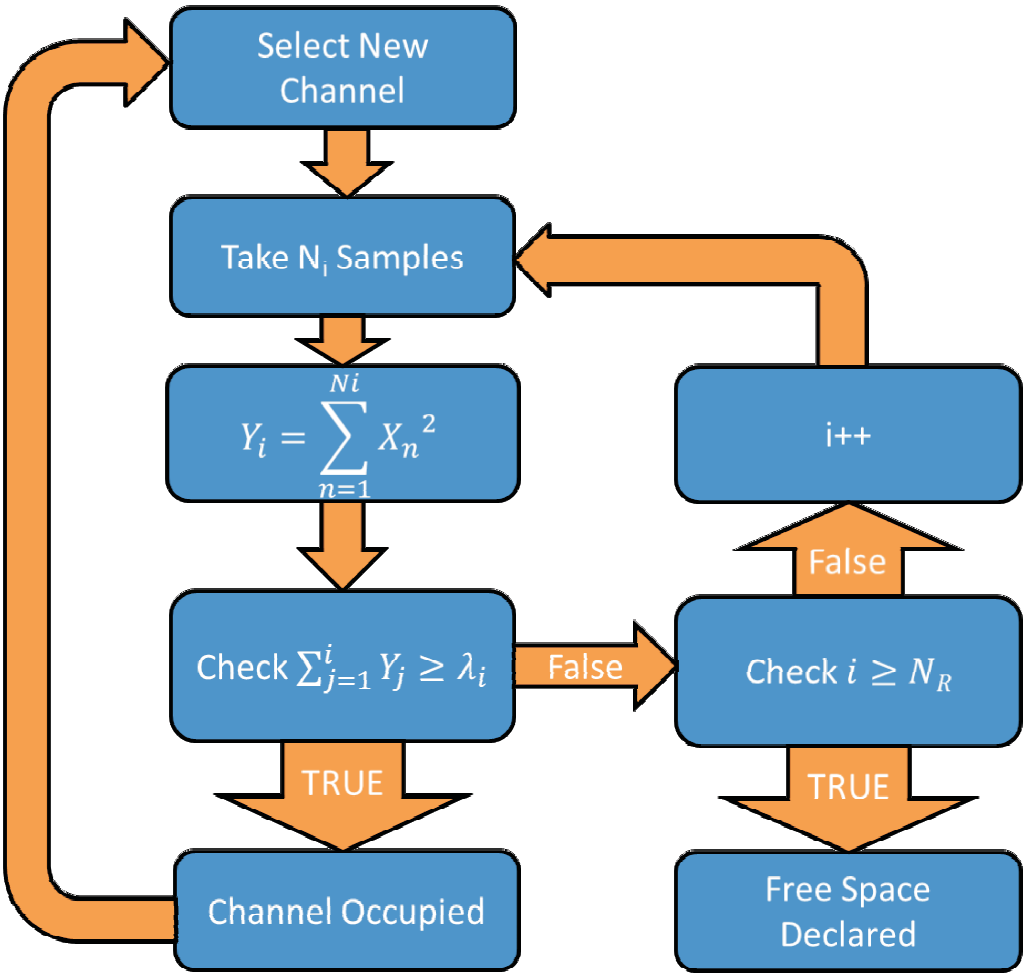

Figure 6.1: Flow diagram for CDFD Architecture

Since the detector makes decisions in the coarse detector mode it is termed a CDFD. This CDFD technique can significantly decrease overall sensing time, as will be shown

in this chapter. However, the system is complicated, having a large number of parameters to choose. For N_R regions there are N_R thresholds to be set, one for each region. The number of samples for each region can also be set. Thus, for a N_R region detector, there are $2N_R - 1$ parameters that can be set; the last detection attempt's threshold and number of samples are constrained to meet the overall P_{fa} and P_{md} rates. Finally, the number of regions can also be changed, adding a further layer of complexity to the calculations.

In this chapter the CDFD architecture is investigated and the characteristic equation derived for the system. A simplified version of the characteristic equation of the detector is derived, which allows a significant reduction in the computational complexity when an optimisation is performed, though at the cost of performance. Using observations on the model, a simplified optimisation is proposed for the full characteristic equation which reduces implementation cost significantly, whilst being sufficiently accurate and show excellent performance. The optimisation results are compared to Monte-Carlo simulations to verify the accuracy of the model. Finally, fading is considered and the impact of fading on the optimisation process is shown.

6.2 Average Number of Samples Required for Detection

To optimise the system (i.e. minimise its sensing time), initially a closed form expression for the average number of samples required by the system must be developed. There are two possible conditions for the channel, namely occupied or unoccupied. Each condition is now considered separately.

6.2.1 Unoccupied Channel

For an unoccupied channel the average number of samples required for detection is a function of the false alarm rate and number of samples per detection attempt. The n^{th} detection attempt $(n = 1, \ldots, N_R - 1)$ has a probability of false alarm of $P_{fa}[n]$, for $N_s[n]$ extra samples (relative to the $(n-1)^{th}$ attempt), where N_R is the total number of regions.

If this detection attempt is reached then $N_s[n+1]$ samples will be used. The average number of samples contributed by this attempt is, therefore, $N_s[n]$ times the probability of a false alarm occurring during the $(n-1)^{th}$ detection attempt. This latter probability is the product of $1-P_{fa}[i]$ for the previous $n-1$ attempts. The average number of additional samples contributed by the n^{th} detection attempt, $N_{unocc}[n]$, is therefore:

$$N_{unocc}[n] = N_s[n] \prod_{i=1}^{n-1} (1 - P_{fa}[i]). \tag{6.1}$$

where $N_s[n]$ is the number of extra samples used by the n^{th} detection attempt, relative to the $(n-1)^{th}$ attempt, and $P_{fa}[i]$ is the probability of false alarm for the i^{th} detection attempt (which can be found using equation 3.15).

The total average number of samples required when analysing an unoccupied channel, N_{unocc}, is the sum of the contribution from each of the N_R regions:

$$N_{unocc} = \sum_{n=1}^{N_R} N_{unocc}[n] \tag{6.2}$$

and, using (6.1), this becomes:

$$N_{unocc} = \sum_{n=1}^{N_R} N_s[n] \prod_{i=1}^{n-1} (1 - P_{fa}[i]). \tag{6.3}$$

This is then the average number of samples required for detection when the channel is unoccupied.

6.2.2 Occupied Channel

When a signal is present with an $SNR \geq -21dB$ the detector is attempting to declare the channel occupied. Each subsequent detection attempt is more likely than the previous one to detect a lower SNR value. Thus, it can be imagined that the PDF is being partitioned on a per-detection attempt basis. This is shown in Fig. 6.2 where the sections of the SNR distribution that each of the detectors detects signals for are shown.

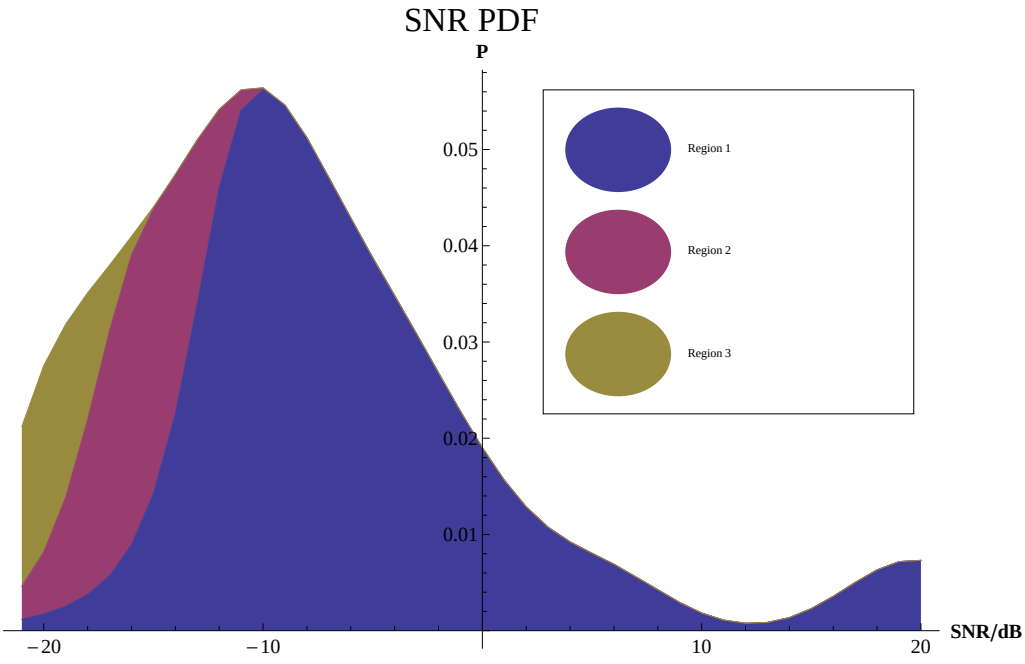

Figure 6.2: Contributions of Each Region for CDFD Spectrum Sensing

The results shown are for a three region detector with the sample distribution shown in Fig. 4.4. The detector values have not been optimised, as this is for illustrative purposes

only. The number of samples used for the first to last detectors, respectively, are 5000, 45000 and 355000 samples. Each detector has a probability of false alarm of 0.02.

The first detection attempt uses the fewest samples, but covers the majority of the distribution. This is to be expected as signals with an SNR greater than -10dB can be reliably detected by an energy detector with 5000 samples. In practice, the last detection attempt does not detect a significant percentage of the signals, but it is required to guarantee that signals at an SNR of -21dB are detected with the required P_{md}.

The average total number of samples required for detection for an occupied channel, N_{occ}, is given by the sum of the average number of samples required for each region.

Hence:

$$N_{occ} = \sum_{n=1}^{N_R} N_{occ}[n], \tag{6.4}$$

where N_R is the total number of regions used and $N_{occ}[n]$ is the average number of samples for the n^{th} region.

Each individual application of the detector analyses a section of the SNR PDF. The contribution of the n^{th} region, $N_{occ}[n]$, is given by:

$$N_{occ}[n] = N_s[n]P_{dr}[n], \tag{6.5}$$

where $N_s(n)$ is the number of extra samples for the n^{th} region and $P_{dr}[n]$ is the probability of the signal being detected during the analysis of the n^{th} region (which can be found using equation 3.13).

A signal with an SNR of γ will only be analysed in the n^{th} region if a missed detection occurs for all of the preceding detection attempts. The probability of this occurring i.e. the probability of reaching attempt n, $P_{dr}(n, \gamma)$, is the product of the missed detection

probabilities for all the preceding detection attempts. This is given by:

$$P_{dr}(n,\gamma) = \prod_{i=1}^{n-1} P_{md}(i,\gamma), \tag{6.6}$$

where $P_{md}(i,\gamma)$ is the probability of missed detection for the i^{th} detection attempt with an SNR of γ.

Weighting equation (6.6) by the value of the SNR PDF at γ, $f_{SNR}(\gamma)$, and integrating over the range of SNR values gives the average probability of detecting at the n^{th} attempt, $P_{dr}[n]$, or:

$$P_{dr}[n] = \int_0^\infty f_{SNR}(\gamma) \prod_{i=1}^{n-1} P_{md,i}(\gamma)\,\mathrm{d}\gamma, \tag{6.7}$$

where f_{SNR} is the user defined SNR PDF whose method of generation is described in section 4.3.3.

Using (6.5) and (6.7), (6.4) becomes,

$$N_{occ} = \sum_{n=1}^{N_R} N_s[n] \int_0^\infty f_{SNR}(\gamma) \prod_{i=1}^{n-1} P_{md,i}(\gamma)\,\mathrm{d}\gamma, \tag{6.8}$$

where $N_s[n]$ is the number of extra samples for the n^{th} region, f_{SNR} is the user defined SNR PDF and $P_{md,i}(\gamma)$ is the probability of missed detection for the i^{th} region at an SNR of γ. However, since it is more appropriate to describe the SNR in the dB scale, the limits of integration change to:

$$N_{occ} = \sum_{n=1}^{N_R} N_s(n) \int_{-\infty}^\infty f_{SNR_{dB}}(\gamma) \prod_{i=1}^{n-1} P_{md,i,dB}(\gamma)\,\mathrm{d}\gamma. \tag{6.9}$$

where $P_{md,i,dB}(\gamma)$ is $P_{md,i}(\gamma)$ for the corresponding dB value of SNR.

6.2.3 Total Average Number of Samples Required

For an occupancy θ, where $0 \leq \theta \leq 1$, the average total number of samples required is:

$$N_{T,CDFD} = \theta N_{occ} + (1-\theta) N_{unocc}. \tag{6.10}$$

Using equations (6.3) and equations (6.8), (6.10) becomes:

$$\begin{aligned} N_{T,CDFD}(\theta) &= \theta \sum_{n=1}^{N_R} N_s[n] \int_{-\infty}^{\infty} f_{SNR_{dB}}(\gamma) \prod_{i=1}^{n-1} P_{md,i,dB}(\gamma)\, \mathrm{d}\gamma \\ &\quad + (1-\theta) \sum_{n=1}^{N_R-1} N_s[n] \prod_{i=1}^{n-1} (1 - P_{fa}[i]) \end{aligned} \tag{6.11}$$

or, for a sampled SNR PDF, in this case satisfying $-21dB \leq \gamma \leq 10dB$, with a step size of Δ:

$$\begin{aligned} N_{T,CDFD}(\theta) &= \theta \sum_{n=1}^{N_R} N_s[n] \Delta \sum_{\gamma=-21}^{10} f_{SNR_{dB}}[\gamma] \prod_{i=1}^{n-1} P_{md,i}[\gamma] \\ &\quad + (1-\theta) \sum_{n=1}^{N_R-1} N_s[n] \prod_{i=1}^{n-1} (1 - P_{fa}[i]). \end{aligned} \tag{6.12}$$

As in equation (5.3), the SNR summation is not strictly a summation, as the step size is not unity. To find the minimum number of samples required for this system would require an optimisation over N_R and all the different values of $N_s[n]$ and $P_{fa}[n]$. Since all of the dimensions are interdependent, solving for a closed form solution is deemed out of scope here.

Added to this difficulty is the fact that $f_{SNR_{dB}}$ is user defined and cannot be assumed to be

of any particular type. Thus, either numerical methods are required to find the optimum multi-resolution detector parameters or significant simplification is required.

The accuracy of equation (6.12) for $N_R = 2$ is shown in Fig. 6.3, where the accuracy of the theory is compared to Monte-Carlo simulations with 50,000 runs. The simulations, which were performed on the same machine as that used to generate the theoretical curves, took approximately 4,000 times longer to generate the data required. The time taken was approximately forty hours for the simulations vs thirty seconds for the analytical model. In addition, the Monte-Carlo curve is not smooth, indicating that more runs would be required to reduce the variability of the relative error. The mean relative error was less than 0.005 (0.5%) and the maximum relative error was 0.015 (1.5%). The plot seems random in its distribution of peaks. This further indicates that more samples were required for the simulations and that the analytical equations predict performance accurately.

This analysis was only done for $N_R = 2$. Visualising similar data for higher dimensions would be difficult and adding another dimension would increase the run time of the simulations beyond that which is practical here. If $N_R = 3$ was used then, with another thirty points in that dimension, the simulations would take approximately ten weeks, whilst, for $N_R = 4$, years would be required, with available hardware. Instead, for higher order systems, the predicted optimum point is tested, i.e. Monte-Carlo simulations are performed for the parameters chosen by the optimisation and the results compared to the expected values from the theory.

6.3 Non-Idealities

Here, two non-idealities are investigated. The effects of sample reuse and time-varying signals on the detector are modelled. It is shown that both have minimal impact and can

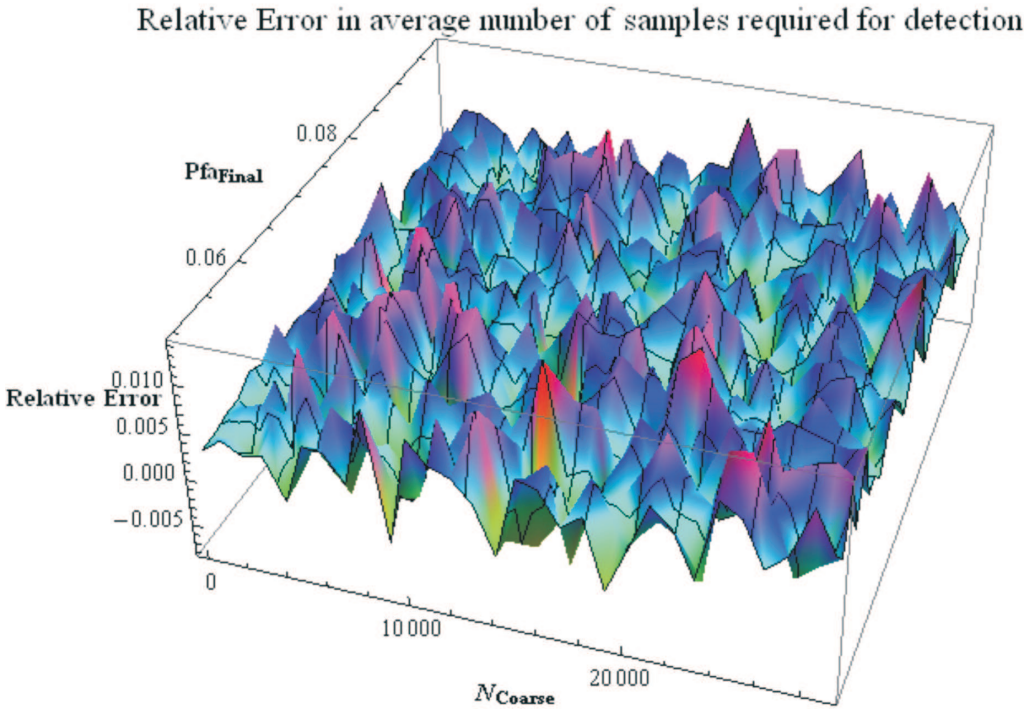

Figure 6.3: Relative Error in Results Compared to Simulations

be safely discounted under normal operating conditions.

6.3.1 Sample Reuse

Sample reuse changes the distribution of the output of the energy detector but allows even greater speed gains to be realised. Even if the samples cannot be satisfactorily used to determine the occupancy of the band, they contain information that may be usefully exploited.

6.3.1.1 No Signal Present

When there is no signal present, an energy detector output Y has a chi-square distribution with N_s degrees of freedom, where N_s is the number of samples used, i.e.:

$$Y \sim \chi^2_{N_s}. \tag{6.13}$$

If the samples from previous detection attempts are used then the total value of the energy detector output at the n^{th} detection attempt, $Y_{T,n}$, is:

$$Y_{T,n} = Y_n + Y_{T,n-1}, \tag{6.14}$$

where Y_n is the contribution due to the extra samples for the n^{th} region and $Y_{T,n-1}$ is the total result for the previous region.

Y_n is a chi-square variable with $N_s[n]$ degrees of freedom or:

$$Y_n \sim \chi^2_{N_s[n]}, \tag{6.15}$$

where, as before, $N_s[n]$ denotes the number of samples used.

The contribution due to the previously acquired samples, $Y_{T,n-1}$, is not a chi squared variable due to the fact that it might not be used if it exceeds the threshold for the $(n-1)^{th}$ region and a signal is declared present.

Instead, it is a censored chi-square variable (i.e. a chi-square variable where any result over the $(n-1)^{th}$ threshold, $\lambda[n-1]$, is not included) or:

$$Y_{T,n-1} \sim \chi^2_{N_{T,n}[n-1]}|_0^{\lambda[n-1]}, \tag{6.16}$$

where $N_{T,n}[n-1]$ is the total number of samples for the $(n-1)^{th}$ region, or the sum of all the values of $N_s[i]$ for $0 \leq i \leq n-1$ and $Z|_a^b$ denotes that the distribution is censored for outputs $z \leq a$ and $z \geq b$, where z is a random variable governed by the pdf Z.

As the ratio of $N_T[n-1]$ to $N_s[n]$ increases, the distortion introduced by the censored

distribution increases. Thus, for lower N_R, the effect is lessened. As P_{fa} decreases, λ increases and, thus, the distortion introduced by the censored distribution also decreases, as less of the distribution is removed by the censoring.

Shown in Fig. 6.4 is the comparison between the expected 'uncorrected' P_{fa} and the actual P_{fa} when the censored distribution is included. All the probabilities of false alarm for the regions are set to the same value $P_{fa,set}$.

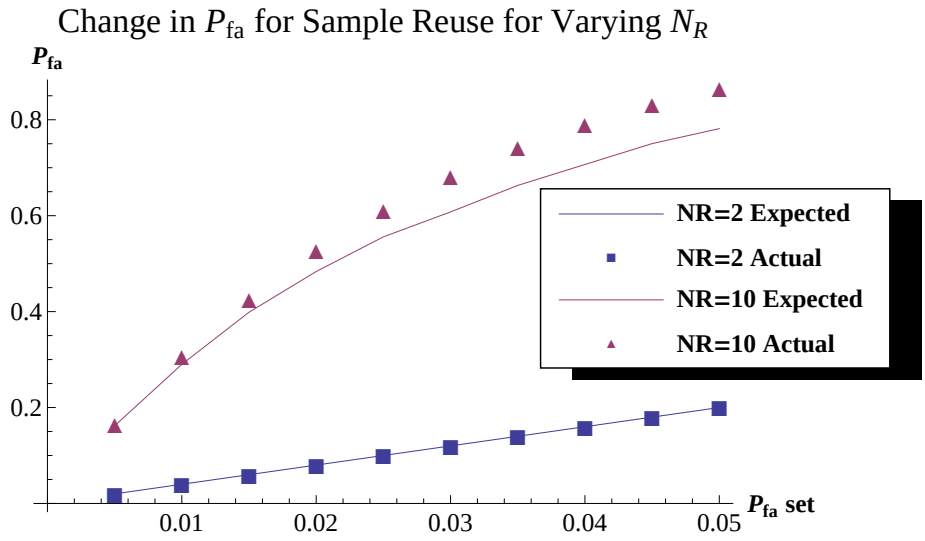

Figure 6.4: Change in P_{fa} that Results from Sample Reuse

It can be seen that for $N_R = 2$ the expected value matches the actual value very closely and there is little distortion. For $N_R = 10$ at a reasonable P_{fa} i.e. the total false alarm rate $P_{fa} \approx 0.16$, the distortion is lessened to an acceptable level (2.6%), at higher levels the distortion grows significantly and the censoring would have to be included in the model. As the IEEE 802.22 standard [3] requires a total false alarm rate of $\leq 10\%$, the overall distortion will be even less for an implementation. It is, therefore, deemed acceptable here to reuse samples without having to model the change in P_{fa}.

6.3.1.2 Signal Present

Similarly, it is possible to show that, when a signal is present, sample re-use does not require modelling for $N_R < 10$. The analysis is complicated by the SNR distribution but, due to the large value of β, any SNR for which there is a non-negligible value of P_d (thus censoring the contribution to the second detector) for the first detector implies there is also a negligible value of P_{md} for the second detector.

Since the system detection probabilities are defined for a signal with an SNR of -21dB, the change in P_{md}, due to the extra attempts, will be negligible for small N_R ($N_R \leq 10$) and will always reduce P_{md}. For $N_R = 10$, the overall P_{md} is still approximately 8% when set to 10%, though this will change with differing values for $N_s[n]$.

This analysis was done for equal P_{fa}. As will be shown later in Section 6.4.3, the optimum solution for the P_{fa} values is heavily weighted towards the last detection attempt. The other values for P_{fa} will be significantly lower. This further reduces the effect of sample re-use as there will be less of a censoring action.

6.3.2 Timing Issues

In section 3.7, and in some of the work studied in Chapter 2 [43, 45], the performance degradation caused by time varying occupancies is considered. Using equation (3.40), the probability of change was found for sensing times, t_{sense} of 0.1s, 0.15s and 0.2s. This corresponds to 100k, 150k and 200k samples, respectively, at a sample rate of 1MSps. The average of the "ON" and "OFF" times, $T_{ON/OFF}$, is varied and the probability of change was found. The result is shown below in Fig. 6.5.

For channels that vary less often than every ten seconds there is a small value of P_{change} of the order of $\approx 1.5\%$ for $T_{ON/OFF}$ =10s and t_{sense} =0.15s. This will not affect performance

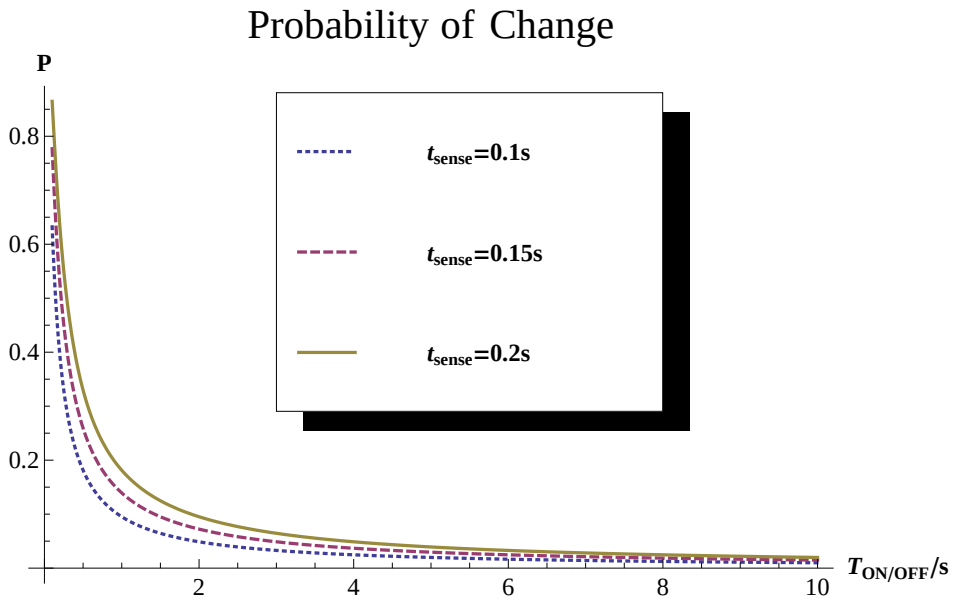

Figure 6.5: Change in P_{change} with $T_{ON/OFF}$ for Varying t_{sense}

significantly. However, for very quickly changing channels ($\leq$ 1s) the value of $P_{change} \approx 14\%$ for $T_{ON/OFF}$ = 1s and t_{sense} =0.15s. If the detector is required to run in such rapidly varying channels, further work is required to ensure acceptable performance, though this is beyond the scope of this work.

If the weighted sum proposed in [43] were to be used with this detector architecture it is possible that the impact of changing occupancies in the channel would be reduced even further, though this was not investigated here. However, it is worth noting that the utility of such a quickly changing channel is unlikely to be high, in general. To reduce the probability of negative impact, the CR would have to scan the channel every second to regenerate the occupancy estimate. This regularity would, in practice, cause substantial power consumption. It is, therefore, considered reasonable that the CR would avoid channels with such rapidly changing occupancies and the detector studied here is

still important.

6.4 Optimisation Options

There are three options for optimising the system based on (6.11).

1. All false alarm rates and sample numbers are variable. Each region has two parameters that are optimised, except for the final region which has only one parameter, namely the false alarm rate.

2. All sample numbers are variable and all false alarm rates are equal. The P_{fa} for each region is also set to be equal and only the number of samples for each region is varied.

3. All sample numbers are variable and the last region's false alarm rate is variable. All other false alarm rates are set equal.

Note that, for all three methods, the system must satisfy the P_{fa} and P_{md} criteria at the SNR specified, i.e. $P_{fa} \leq 0.1$ and $P_{md} \leq 0.1$ at an SNR of -21dB. The final region is used to guarantee the missed detection probability; thus, the final region is constrained and only one of the parameters, namely P_{fa} or N_s, can be varied freely.

Of critical importance when using the CDFD is ensuring that the overall P_{fa} and P_{md} specifications are not exceeded. Compared to a single detection attempt with the same P_{fa} and P_{md} as the last detection region, multiple detection attempts will reduce the overall P_{md} but increase the P_{fa}.

To guarantee an overall P_{fa} satisfying $P_{fa} \leq P_{fareq}$, for example $P_{fareq} = 0.1$, under N_R detection attempts, the system requires setting the individual values for $P_{fa}[n]$ such that:

$$P_{fareq} = 1 - \prod_{n=1}^{N_R} 1 - P_{fa}[n]. \tag{6.17}$$

This constraint removes a degree of freedom and reduces the total variables to be optimised from $2N_R$ (all the thresholds and numbers of samples) to $2N_R - 1$ (where, generally, the number of samples used in the last attempt is set to guarantee performance at -21dB).

6.4.1 All False Alarm Rates and Sample Numbers Variable

In this case no simplifications are made when attempting to optimise, all the parameters are variable. Equation (6.11) is used without modification and, using numerical optimisation techniques, the optimum parameters can be found for this system such that N_T is minimised.

6.4.2 Equal Probabilities of False Alarm

If the optimisation is simplified such that all values of $P_{fa}[n]$ are set to be equal then (6.17) can be re-arranged to become:

$$P_{fa}[n] = 1 - (1 - P_{fareq})^{\frac{1}{N_R}}. \tag{6.18}$$

For a system with $N_R = 4$ and $P_{fareq} = 10\%$ the individual values for $P_{fa}[n]$ are set equal to 0.026 for $1 \leq n \leq N_R$.

For a detector with the requirements $P_{fa} \leq 10\%$ and $P_{md} \leq 10\%$ at an SNR of -21dB, if no signal is present then the detector requires an increasing number of samples as N_R increases. This is shown in Fig. 6.6. A single region energy detector requires 209k samples under the same conditions. Optimisation will not greatly change the number of

samples required when the channel is unoccupied, if the individual values for $P_{fa}[n]$ are set to be equal.

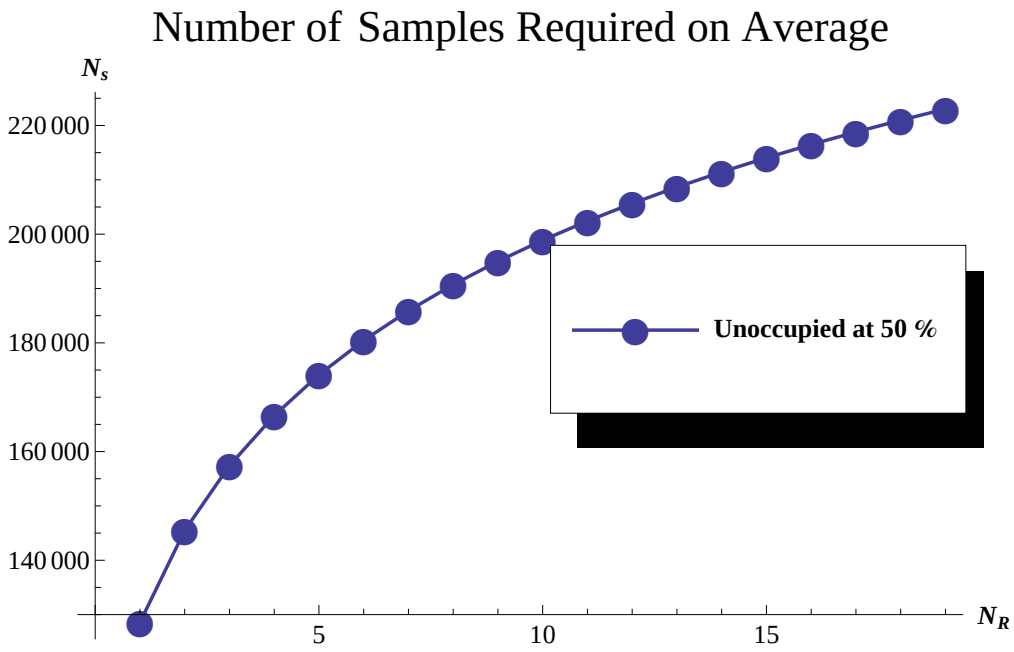

Figure 6.6: Increase in the Number of Samples Needed for $P_{fa} \leq 10\%$ and $P_{md} \leq 10\%$ at -21dB, Without a Signal Present, for Multiple Regions

When optimising the system it must be noted that this increase will counteract the speed gain that extra regions will provide when detecting signals. For $N_R \geq 13$ there is no gain possible, as the average number of samples required for the unoccupied case at 50% occupancy is greater than the average number of samples required for naive detection.

6.4.2.1 Number of Regions

For an unknown environment, the average number of samples required can be approximated by the average over all the occupancies.

$$N_T = \int_0^1 N_T(\theta)\mathrm{d}\theta. \tag{6.19}$$

i.e.

$$N_T = \int_0^1 N_{unocc}(1-\theta) + N_{occ}\theta. \tag{6.20}$$

Since the number of samples does not depend on the occupancy, this gives $\theta = 0.5$.

As N_R increases N_{unocc} will increase also, as can be seen in Fig. 6.6.

Other distributions or occupancies may have other optimum values for N_R but here this work is concerned with distributions that are based on a lack of knowledge of the system (uniform distribution, 50% occupancy) or ones based on real systems previously investigated [106]. For the distributions used here $N_R = 2$ is optimum when $\theta = 0.5$. This is shown in Fig. 6.7. It can be seen in Fig. 6.7 that, at 75% occupancy, N_R =3 gives the maximum possible speed gain. With this occupancy, however, the difference between the N_R =2 and N_R =3 detectors is small.

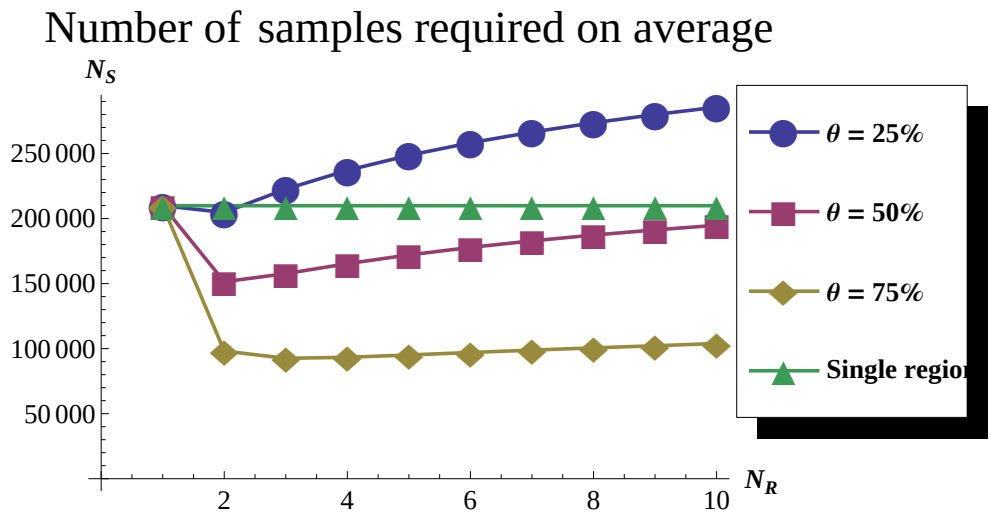

Figure 6.7: Average Number of Samples Required vs N_R for Varying Occupancies

Thus, it is reasonable to make the assumption that, for the distribution adopted here, $N_R = 2$ is the optimum for the case of P_{fa} being equal over all detection attempts. For other distributions this is not guaranteed but, for similar distributions, it is expected that $N_R = 2$ provides either the optimum, or a result sufficiently close to the optimum for practical purposes. This method is concerned with reducing the optimisation problem, thus $N_R = 2$ is more attractive as it reduces the number of parameters being optimised to one, namely the number of coarse detector samples.

With $N_R = 2$, as derived for the simplified case, where all values for P_{fa} are equal, equation (6.9) can be greatly simplified and becomes:

$$N_{occ} = N_c + (N_f - N_c) \int_{-\infty}^{\infty} f_{SNR_{dB}}(\gamma) P_{md_c}(\gamma)\, \mathrm{d}\gamma, \tag{6.21}$$

where N_f is the number of samples required to guarantee performance at an SNR of -21dB, $P_{md_c}(\gamma)$ is the probability of missed detection for the coarse detection attempt and N_c is the coarse region number of samples that the system is optimising.

Since, for practical applications, $f_{SNR_{dB}}$ is a sampled distribution, the integral in (6.21) is replaced with a sum to become:

$$N_{occ} = N_c + \Delta(N_f - N_c) \sum_{\gamma=-21dB}^{10dB} f_{SNR_{dB}}[\gamma] P_{md}[\gamma], \tag{6.22}$$

where Δ in the summation depends on the smoothness of the $f_{SNR_{dB}}$ and, here, $\Delta = 0.5dB$.

6.4.3 Last Region's False Alarm Rate Variable

The most important value for P_{fa} is the final value, $P_{fa}[N_R]$. To reduce the computational complexity of finding the optimum values for the system, the number of variables being optimised can be reduced. If $P_{fa}[N_R]$ is optimised, and the rest of the $P_{fa}[n]$, $1 \leq n \leq N_R - 1$, are set equal then, for N_R regions, the optimisation requires optimising over N_R variables instead of over $2N_R - 1$ variables.

The individual values $P_{fa}[n]$, $1 \leq n \leq N_R - 1$, are given by:

$$P_{fa}[n] = 1 - \frac{(1 - P_{fareq})^{\frac{1}{N_R - 1}}}{1 - P_{fa}[N_R]}. \tag{6.23}$$

In the next section, an optimisation is performed for each of the three options and the methods are compared. It will be shown that, for the test distribution considered here, the last method, only the last false alarm rate being variable, is the most appropriate.

6.5 Optimisation of CDFD

The simplified equation, (6.12) and the full complexity equation, (6.22), are now optimised. The distributions considered here are the distributions based on the Irish DTV system, shown in section 4.4, and later, in Chapter 7, some Gaussian distributions with similar shape but higher SNR and a uniform distribution (representing a complete lack of prior knowledge of the SNR distribution, where all SNRs are equally likely in the range -21dB to 0dB) will be considered.

The curves generated for the number of samples are smooth and, in the cases studied here, there is only one minimum, the global minimum. Therefore, numerical optimisation is possible and can use simple optimisation schemes such as Pattern Search [115] or Nelder-

Mead [116]. If the distribution is more complicated and/or multimodal, then the curves may not have a single minimum and a more sophisticated optimisation scheme may be needed. In this work, the Pattern Search algorithm is used as it is easily implemented and relatively efficient at finding the maximum or minimum of a numerically computable function. The Pattern Search method is discussed further in Section 6.7.1.

Instead of using equation (6.22), the exact equation, equation (6.12) is used for the optimisation, where all values of P_{fa} are set to be equal for equal regions. This allows a comparison between all the optimisation options over a range of N_R values. As would be expected, the lowest average number of samples occurs for N_R = 2 in this case. The most complicated optimisation allows the independent setting of P_{fa} for all the regions. Finally, the system is optimised when only the last P_{fa} is variable.

Using numerical optimisation methods, the optimum values for the test distribution are found for a range of N_R values. The minimum number of samples required was found for each distribution and the results are plotted in Fig. 6.8. Note that the exact and approximated optimisation results overlap until N_R = 5 and the symbols for the approximated optimisation simulation results are masked by the symbols for the exact optimisation simulation results.

In addition, Monte-Carlo simulations, with 100,000 runs, were performed to verify the results of the optimisation for each value of N_R. In most cases, the predicted value of N_T is shown to be within 1% of the Monte-Carlo values. The only cases where the predicted values do not match the Monte-Carlo simulations well are for large N_R in the simplified case. This is partly due to the fact that the assumptions about sample reuse, stated in Section 6.3.1, are not as valid for such large values of P_{fa} in the earlier detection attempts.

There is almost no degradation in performance when switching to the less complex op-

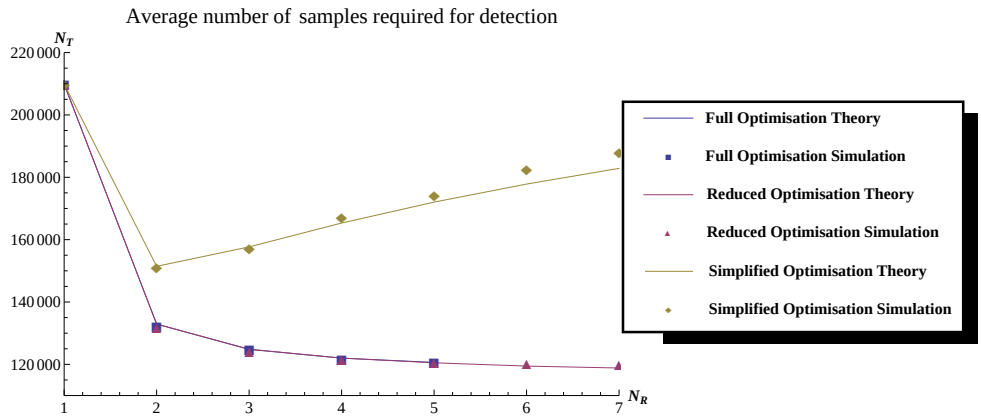

Figure 6.8: Optimised Speed Gain Compared with Monte-Carlo Simulations

timisation where only the last P_{fa} is variable. The maximum percentage difference was found to be 0.1%, when $N_R = 5$. Further exact optimisation beyond $N_R = 5$ was found to be difficult, even with the in-built Mathematica® 8.0.1.0 optimisation algorithms, due to the number of variables and the constraint on total P_{fa} . The simplified method does not suffer from this problem, and the optimisation evaluates significantly faster. Since the accuracy is very good and the reduction in computational complexity is so significant, this simplified optimisation will be used for the rest of this work.

6.6 CDFD in Fading Channels

An important section of the IEEE 802.22 standard is the requirement for the system to maintain satisfactory performance even when operating in a fading environment. If a signal is subjected to fading then the SNR, γ, of the signal is distributed according to the fading distribution. For example, a signal with an average SNR of $\overline{\gamma}$, when subjected to Rayleigh fading, has its SNR γ distribution according to [55]:

$$f(\gamma) = \frac{1}{\overline{\gamma}} exp(-\frac{\gamma}{\overline{\gamma}}), \gamma \geq 0. \tag{6.24}$$

Similarly when there are multiple fading paths Nakagami Fading is used [55], and,

$$f(\gamma) = \frac{1}{\Gamma(m)} (\frac{m}{\overline{\gamma}})^m \gamma^{m-1} exp(-\frac{m\gamma}{\overline{\gamma}}), \gamma \geq 0, \tag{6.25}$$

where m is the Nakagami parameter, and $\Gamma(m)$ is the gamma distribution evaluated at m and m is the measure of the multipath effects.

Applying fading to the original SNR distribution creates a new distribution. This distribution can then be used in place of the original SNR distribution to find the optimum number of samples for the coarse detector. This was covered in detail in section 5.8.

The number of samples required by the fine detector needs to be determined for 10% false alarm and 10% missed detection at an average SNR, $\overline{\gamma}$, of -21dB; thus, the value for N_{fine} will increase relative to the non-faded case.

Under Rayleigh fading and $P_{fa} = 0.1$, $N_{fine} \approx$ 4.8M samples, whilst, under Nakagami fading with $m = 2$ and $P_{fa} = 0.1$, $N_{fine} \approx$ 0.99M samples. The value of N_{fine} depends on the value of P_{fa} used for the final detection attempt. It is possible to solve for the value of N_{fine} numerically each time, or a look-up table of values for P_{fa} can be used.

The approximate optimisation is performed on the DTV distribution under Nakagami and Rayleigh fading and the results are shown in Fig. 6.9. Also shown are the results of optimisation for a Gaussian channel. In addition, Monte-Carlo simulations with 100,000 runs were performed to verify the result of the optimisation for each value of N_R under the two fading types. In most cases the predicted value of N_T is shown to be within 0.1% of the Monte-Carlo values.

It is difficult to see the slope of the curves in Fig. 6.9 due to the large difference in the

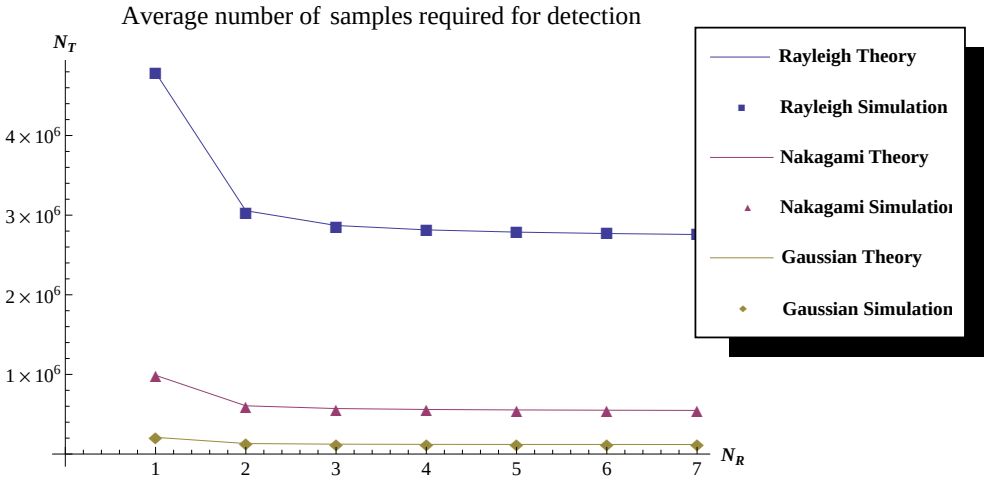

Figure 6.9: The change in N_T with Optimised Sensing Under Different Fading Types for Approximate Optimisation

number of samples required under the different fading types. Shown in Fig. 6.10 is the speed gain, relative to a single region detector with P_{fa}=10% and P_{md}=10% at a SNR of -21dB, for the three channels considered here.

As can be seen in Fig. 6.10, the gain for adding extra regions decreases with each region added. The cost of optimisation increases with every additional region, e.g. with the pattern search method the number of test points per iteration doubles with every extra dimension added. For the test distribution used here, the optimum values for the system are shown in Table 6.1.

For a system in Rayleigh Fading with the same underlying distribution, Table 6.2 is produced.

This analysis assumes complete knowledge of the SNR distribution. Naturally, in a real system, there will almost certainly be an error in the estimated distribution relative to the actual distribution. This will now be investigated.

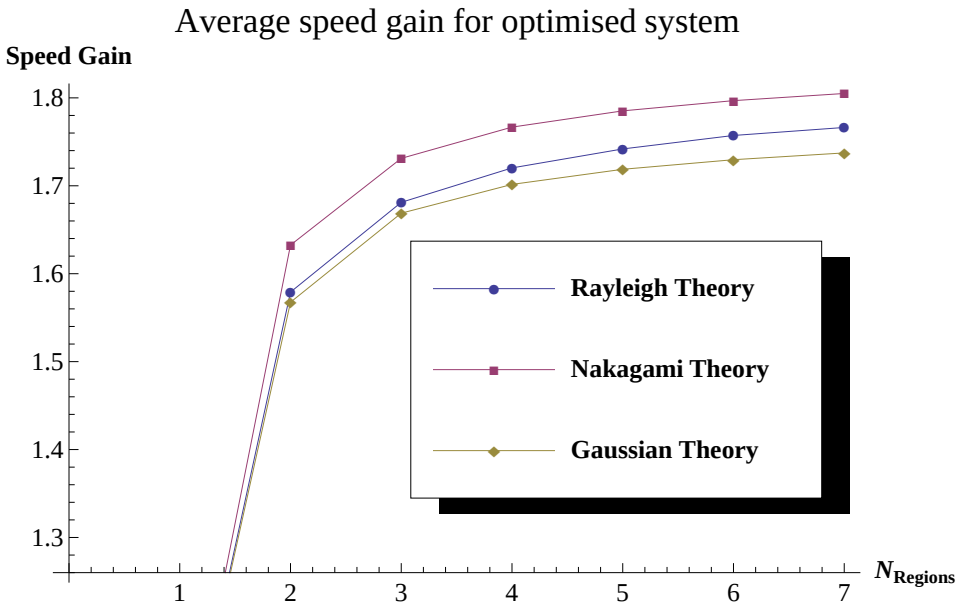

Figure 6.10: The change in Speed Gain with Optimised Sensing Under Different Fading Types for Simplified Equation

Table 6.1: Comparison of Optimum Points for CDFD for Gaussian Channel

Gaussian	N_R=1	N_R=2	N_R=3	N_R=4	N_R=5	N_R=6	N_R=7
N_1	209k	19.6k	7.7k	4.6k	3.28k	2.5k	2.1k
N_2		215k	57.6k	29.3k	18.8k	14.3k	10.9k
N_3			214.5k	82.9k	50k	38.2k	26.8k
N_4				214.2k	97.1k	68.9k	49.1k
N_5					214.1k	102.1k	77.1k
N_6						214.1k	112.3k
N_7							213.9k
P_{fa}	0.1	0.0947	0.0952	0.0955	0.0956	0.0956	0.0958
N_TTheory	209k	133.9k	124.8k	122k	120.5k	119.4k	118.8k
N_T Simulation	209k	132k	124k	121.6k	120.6k	120.4k	119.9k

Table 6.2: Comparison of Optimum Points for CDFD for Rayleigh Channel

Gaussian	N_R=1	N_R=2	N_R=3	N_R=4	N_R=5	N_R=6	N_R=7
N_1	4.79M	396.2k	140.9k	78k	53k	40.1k	32.4k
N_2		4.86M	1.2M	569.7k	353k	251.3k	195k
N_3			4.84M	1.79M	1.03M	694.4k	519.3k
N_4				4.84M	2.16M	1.39M	1M
N_5					4.84M	2.38M	1.67M
N_6						4.84M	2.54M
N_7							4.836M
P_{fa}	0.1	0.0982	0.0986	0.0987	0.0988	0.0988	0.0988
N_TTheory	209k	3.055M	2.87M	2.815M	2.79M	2.77M	2.76M
N_TSimulation	209k	3.042M	2.866M	2.82M	2.8M	2.79M	2.77M

6.7 Optimisation Using a Learned Distribution

In a cognitive radio system the SNR distribution must be learned before the optimisation. This learning was performed in Chapter 4 using KDE. In practice, the optimisation would only be performed periodically, depending on how often the distribution is updated, and could be improved upon by a more efficient optimisation, if necessary. Thus, timing is not of great significance. In addition, rather than using the inbuilt Mathematica® 8.0.1.0 optimisation algorithms, the pattern search algorithm will be used to find the minimum for the estimated SNR distribution. This is because the algorithms in Mathematica® 8.0.1.0 are heavily optimised for the platform and this work is designed for platform independence.

6.7.1 Pattern Search

The pattern search algorithm is a relatively simple search scheme for numerical optimisation. It is a heuristic scheme, with initial settings generally being application specific. A simple implementation of the pattern search algorithm is used here to minimise the

objective function.

Let x_i denote the i^{th} dimension of freedom, or i^{th} variable of the object function being minimised. The function is evaluated for $x_i \pm \Delta x_i$, where Δx_i is the step size per iteration, for x_i. The smallest result is then selected and the process repeated. If there are no smaller results than the current result, the Δx_i are halved for all x_i. The process is then repeated until a specified number of reductions in size are performed. This is illustrated in Fig. 6.11, where the flow diagram is shown.

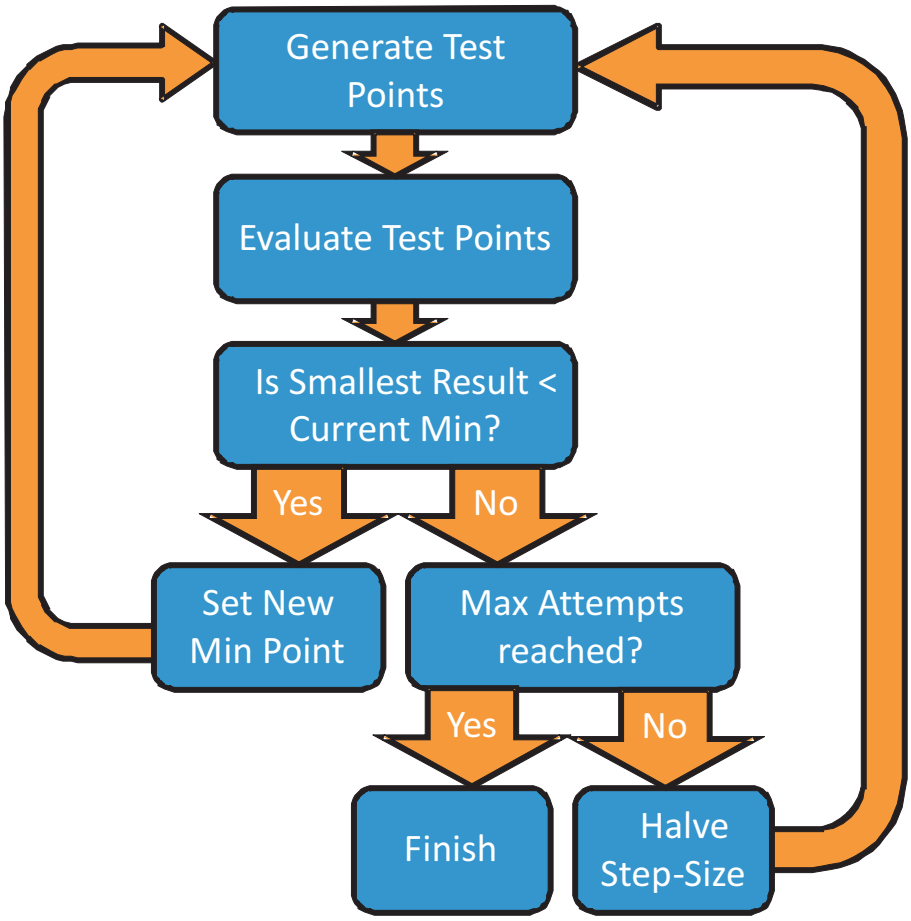

Figure 6.11: Pattern Search Flow Diagram

In general, the computational complexity is impossible to derive in advance as the number

of steps required depends on the object function and the initial values. The number of times the function is calculated per step is twice the number of parameters that can be varied. Thus, for the system considered here, seven regions would require calculating the value of the objective function fourteen times per step. It is also reasonable to assume that, as the number of parameters increases, the number of steps required to find the optimum increases also [117]. Though this assumption cannot be proven here, is has been found to be true, in practice, for the distributions considered here.

6.8 Noise Uncertainty

This system is susceptible to problems due to noise power uncertainty. There are a number of possible solutions to the problem of noise uncertainty.

Firstly, the system could use an accurate noise power estimate. If the noise uncertainty is sufficiently small so as to render its effects negligible (e.g. $\rho \leq 0.001dB$) then the analysis performed here can be used without any alteration.

If the system cannot guarantee a sufficient level of accuracy with the noise power estimate, then another solution must be found. If there are a large number of free bands then the system can sacrifice P_{fa} to reduce the required noise power estimate accuracy. By increasing P_{fa}, the P_{md} decreases and then the overall probability of interference can be kept under the 10% required.

If it is unlikely that there are a number of free channels or that the noise power estimate cannot be reliably found, then co-operative sensing can be used. Co-operative sensing reduces the required accuracy that the system needs to detect a signal at -21dB, in addition the effects of noise uncertainty are reduced. The exact effects of this co-operative mode have not been investigated and cannot be quantified at this time.

Finally, if a single sensor is used without any of the above options then the final option is to use the CSFD with a fine detector that is immune to noise uncertainty such as the CFD.

6.8.1 Performance of Optimised System with Target MISC

An estimate of the SNR distribution is found for a threshold value of $MISC$. This ensures a relatively accurate estimate. Then, using the pattern search algorithm, the system is optimised for N_R, $2 \leq N_R \leq 7$. The values selected by the optimisation are then used with the correct underlying distribution to find the actual performance of the system, under these conditions. This is repeated 1500 times and then averaged to find the new average number of samples required for detection. Finally, the target value of $MISC$ is changed and the process repeated. The results of this analysis are shown in Fig. 6.12.

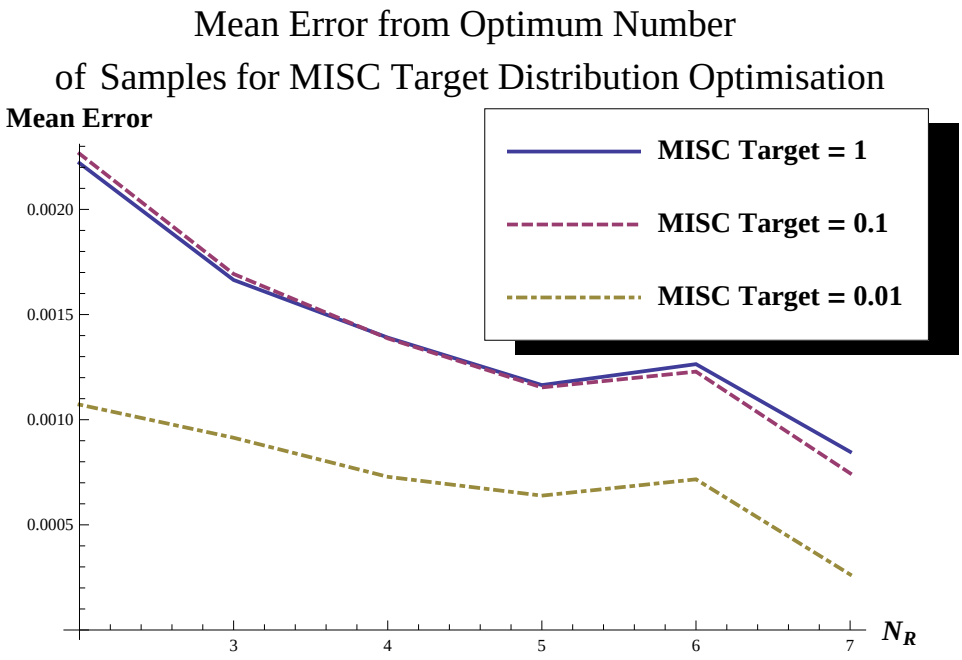

Figure 6.12: Accuracy of CDFD Optimisation vs Target $MISC$ for Estimated SNR Distributions

As can be seen, even a small number of samples can lead to accurate results from the optimisation. This is partly due to the fact that the output is relatively uniform at the maximum and is weakly influenced by small differences in the parameters. For a target $MISC$ of 0.01 and seven regions the mean error is less than 0.1%.

The mean error is small, despite the relatively large target $MISC$. As the $MISC$ decreases, the mean error decreases. This is clearly seen as the target $MISC$ decreases from 0.1 to 0.01. It is less visible as the $MISC$ decreases from 1 to 0.1. This is due to the setup of the $MISC$ simulations. The smallest number of samples that can be used is forty, one set of twenty samples to begin and then another set of twenty samples to compare and generate the $MISC$. The target $MISC$s of 1 and 0.1 do not require, on average, more than this number and, thus, the results are quite similar. If the minimum number of samples was decreased then the results for the $MISC$ target of unity would almost certainly change and the error would increase.

The variance of the error found in the 1500 runs is shown in Fig. 6.13. Similar to the mean error, the variance of the error decreases with the more stringent target value for the $MISC$. Additionally, the difference between the target $MISC$s of 1 and 0.1 is small, though more noticeable in this case. For a target $MISC$ of 0.01 and seven regions the mean error is less than 1.18×10^{-4}%.

The relative lack of accuracy required for the optimised parameters leads to an important question: is optimisation required? Is there a set of values for P_{fa} and N_T that give performance sufficiently close to the optimum across a wide range of distributions? This is investigated in section 7.3.2.

This system has a number of features that make it an attractive solution for cognitive radio sensing schemes. It uses an energy detector to perform the spectrum sensing, thus making it cheap and simple to implement. The optimisation scheme uses only data that is readily

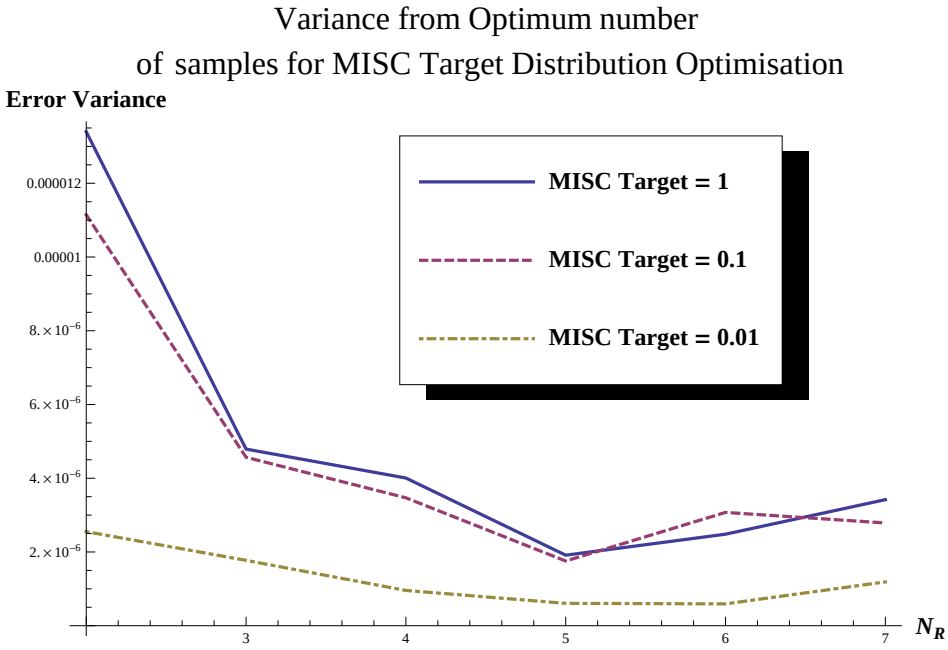

Figure 6.13: Variance of System vs Target $MISC$ for Estimated SNR Distributions

available from the energy detector output, therefore the optimisation scheme is efficient. By using a pattern search the scheme is low on computational complexity. Indeed, the pattern search initial step size can set by using the $MISE$ between the new distribution and the old distribution and the initial point can be the optimum point for the previous distribution. This new, novel, detector architecture is, therefore, a self-optimising scheme that is faster than typical energy detector schemes but also one that does not require any substantial processing.

6.9 Conclusion

In this chapter the CDFD was introduced. By using a multi-region detector, where each consecutive region is analysed with an energy detector of increasing sample size, the

CDFD can be used to reduce the average number of samples required for spectrum sensing in CR applications. By accounting for the SNR distribution, the CDFD can provide significant performance gains over a single region energy detector. This detector, however, does not contain any robustness to noise uncertainty and, therefore, is not effective in situations where the noise uncertainty is expected to be large.

By analysing the CDFD an expression for the average number of samples, equation (6.22), was derived. This expression was shown to be accurate over a range of typical values for the number of regions, number of samples in each detection attempt and the false alarm rate of each detection attempt. The expression also allows user defined SNR distributions to be used. In addition, the average occupancy of the channel can also be included in the calculation, if known. Using Monte-Carlo simulations, for comparative purposes, the average error for a two region detector over a wide range of operating parameters was found to be approximately 0.5%. For larger numbers of regions the result was shown to be accurate at the optimum points also, with the error being less than 0.9% for a seven region detector. This was also shown to be accurate under fading conditions, with a maximum error of 1.5%.

Some of the non-idealities of the system were also investigated. If the samples are reused from one detection attempt to the next, then an error will be introduced into the prediction of the average number of samples. It was shown, however, that for the small number of regions considered here (less than seven, typically) the effect of sample reuse is minimal and can be usefully discounted in the analysis.

Timing issues were also considered. If the occupancy of the channel changes between detection attempts, then, obviously, the detector performance will suffer. For the sensing times considered here, there is little chance of the channel occupancy changing during a detection attempt where the average time between changes is large. For an average

time between occupancy changes of 10s there is approximately a 1.5% chance of the occupancy changing. If the occupancy changes are less frequent, then the probability of change will be even smaller.

After verifying that the expression remained accurate across a range of conditions, multiple schemes for optimising the CDFD were discussed. The nature of the problem indicated that an analytical solution would prove difficult to find and, therefore, numerical optimization methods were used. A simplified (and approximate) optimisation was formulated that reduced the number of variables to be optimised by almost 50% but that achieved performance within 0.1% of the case where exact optimisation was performed. This simplified optimisation requires only the false alarm rate for the final region to be optimised, rather than the false alarm rates for all of the regions. This reduction in the number of variables allows the numerical optimization to be performed more quickly and with less computation, though the exact reduction is difficult to quantify, in general.

In a real implementation the detector would probably not have complete knowledge of the SNR distribution when optimising. Instead, an estimate would have to be used. This estimate is provided using the method in Chapter 4. The method is shown to be very accurate for setting optimum parameters with a mean error of less than 0.1% for a target $MISC$ of 0.01 and seven regions. This new, novel, detector architecture is, therefore, a self-optimising scheme that is faster than typical energy detector schemes but also one that does not require any substantial processing.

This method does not help, however, with the selection of the order in which to scan the channels. If the first region was used to sort a group of channels, then it is likely that the system performance would be increased further over naive detectors. This is investigated in section 7.4 as a Hybrid Coarse-Fine Detector (HCFD) system.

7

Practical Coarse Fine Detectors

7.1 Introduction

In this chapter, a practical CR spectrum sensing architecture will be designed and then tested by simulations and implementation on the IRIS system. Firstly, the requirements for a practical CR system are discussed. A CR system must be realisable, robust, practical and effective to be considered worthwhile. The previous architectures, CSFD and CDFD, are investigated under optimisation for a uniform SNR distribution. This uniform distri-

bution approach is shown to be effective over a wide range of distributions, whilst also removing the requirement for in-line optimisation and SNR distribution estimation.
A new architecture, the HCFD, is introduced and discussed. Using simulations, it is shown to be superior to the CSFD and CDFD in terms of performance, without increasing the implementation cost significantly. Finally, the four architectures, CSFD, CDFD, HCFD and the energy detector are implemented on the IRIS system. It is shown that the architectures can be used in a real environment and that the results match those predicted by simulations to within an acceptable level. The source of the differences between the simulations and practical results will also be discussed.

7.2 Practical Cognitive Radio Receivers

An important feature of the work in this chapter is the implementation of the CF detector on the IRIS system. To date, a large number of possible solutions to the spectrum sensing problem have been proposed and simulated. However, to the author's best knowledge, most of the architectures have not yet been implemented in reality. The implementation of the detectors is of great importance for many reasons. Firstly, while some detector architectures may work well in theory, their real performance may differ significantly. An architecture that is not sufficiently robust to the issues that occur in reality is not, in general, an architecture that can be used for CR applications. In addition, the actual performance of the system will, most likely, not match the predictions sufficiently well. A good model of the system can be used to make useful predictions of the practical performance, whilst implementation testing can also be used to further refine the model.
Some of the criteria for a potential CR system are:

1. It must be realisable, i.e. there cannot be memory/hardware requirements that can-

not be met by current technology.

2. It must be robust, i.e. the radio must be able to operate even in relatively difficult signalling conditions.

3. It must be practical, i.e. the implementation cost of the system must not be too high, either in terms of the resources required or the computational complexity.

4. It must be effective, i.e. there must be no better alternative available.

5. It must be worthwhile, i.e. the system must be superior in some area over other architectures that justifies the implementation cost of the system.

The architectures considered here are realisable. They have been implemented on the IRIS system using USRP front-ends. The architectures do not require any extra hardware to operate when compared to an energy detector.

The architectures can also be considered robust. Even with the USRP front-end, which is not designed for spectrum sensing and has several non-idealities that degrade its performance, in practice the architectures were able to detect signals relatively reliably. The main factor that reduces the effectiveness of any energy detector based architectures is the presence of noise uncertainty. If there is no significant noise uncertainty then the architectures will operate with improved performance.

The architectures, without in-line optimisation, do not require significant extra computation when compared to the energy detector. Thus, practicality is not an issue. Using in-line optimisation for the architectures does require extra computation. If the in-line optimisation can be removed then the architectures become significantly more usable. This is investigated in section 7.3.

It is difficult to declare the architectures considered here "better" than other options. It is possible to say, however, that no superior scheme has been shown, to the author's best knowledge, for CF sensing under an SNR distribution.

Finally, the systems are worthwhile. It will be shown that all the architectures significantly improve upon the energy detector. For example, the HCFD requires less than 30% of the samples required by the energy detector, at a high occupancy ($\theta = 0.9$).

It must be noted that here it is assumed that there is no time penalty for changing channels. This can be achieved by using multiple PLLs to "preload" the next frequency to be looked at.

The optimisation methods shown in previous chapters, Chapter 5 and Chapter 6, are not considered appropriate for the practical implementations studied here. The methods do provide a very accurate estimate of the optimum operating points; however, as will now be shown, there is another solution that provides a set of parameters that are sufficiently close to the optimum for most practical applications. Although other applications may use the optimisation methods of Chapters 5 and 6 to provide increased performance, the cost is deemed too large here and an alternative option is proposed.

7.3 Uniform Distribution

If no estimation of the SNR is performed; the system is unable to make any assumptions about the shape of the SNR PDF. Thus, a possible alternative to estimating the SNR distribution is to assume a uniform distribution, i.e. no knowledge of the system, and to optimise for that distribution instead.

To test the validity of this approach the test distribution from Chapter 4 is not used. Since using only one distribution is not sufficient to test the method, a range of Gaussian dis-

tributions is used instead. Each Gaussian distribution has a mean value between -19dB and -1dB and a standard deviation between 1dB and 15dB (variance between 1dB and 225dB).

7.3.1 CSFD Optimised for a Uniform Distribution

Firstly, the optimisation of the CSFD under a uniform distribution is considered. For a uniform distribution, in the range -21dB to 10dB, the optimum number of coarse samples, N_{coarse}, is 1200 samples.

To see the how well this number of coarse samples performs, consider a Gaussian distribution with a mean of -5dB and a standard deviation of 7dB. For a CSFD optimised with full knowledge of this distribution, the optimum number of coarse samples, N_{coarse}, is 1500 samples, with an average speed gain of 1.81.

Using a uniform distribution, with N_{coarse}= 1200 samples, the average speed gain is 1.8 times. The speed gains for the two methods have only a 0.16% difference.

Repeating this analysis over the range of Gaussian distributions gives the contour plot in Fig. 7.1. The relative speed gain for a CSFD optimised for the uniform distribution compared to a CSFD optimised with full knowledge of the SNR distribution is shown.

The relative speed gain is close to the optimum for a wide range of distributions. The lowest speed gain is 89% of the optimum, for a distribution with a very small variance. The other distributions have significantly higher relative gains. Using the PDF method described in Chapter 5, the simulations were completed in less than two weeks on the boole cluster [118]. Without this model the simulations would have taken 3-4 years on the same cluster.

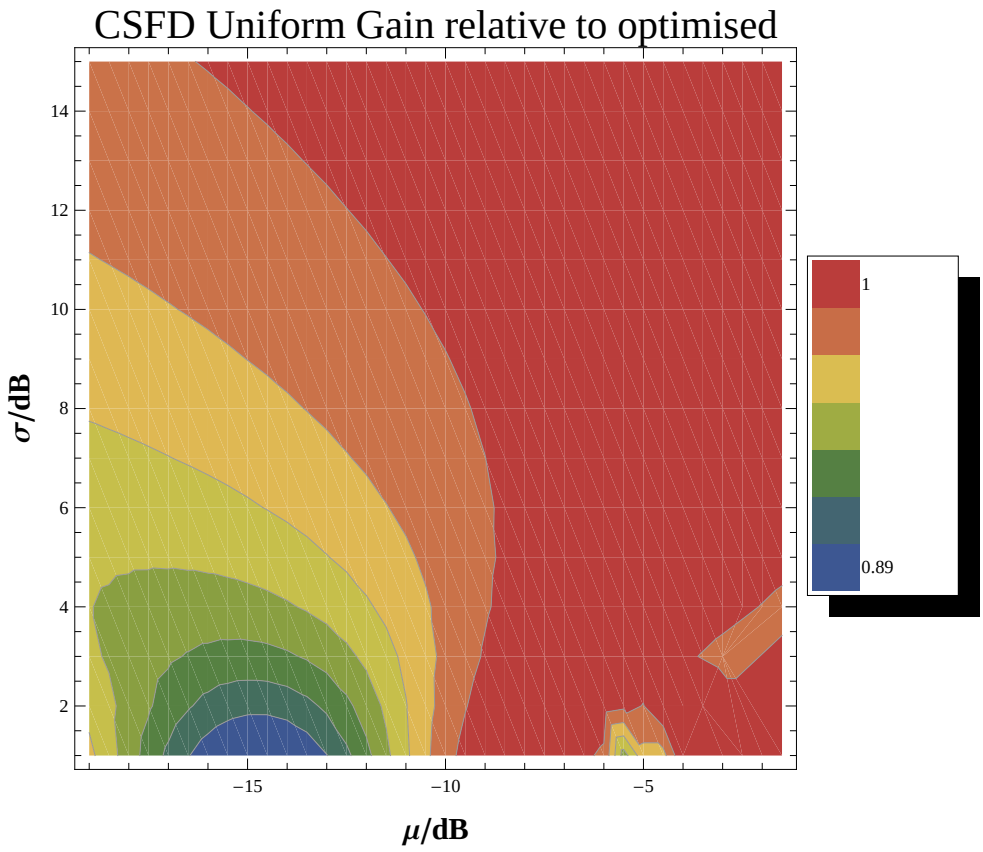

Figure 7.1: Relative Speed Gain for Uniform Distribution

7.3.2 CDFD Optimised for a Uniform Distribution

For example, consider a Gaussian distribution with a mean of -5dB and a standard deviation of 7dB. For a CDFD with N_R = 4, and full knowledge of this distribution, the optimum sample values are $N_{samples}$ = (2000, 17000, 64500, 211540) with P_{fa} = (6.5×10^{-4}, 6.5×10^{-4}, 6.5×10^{-4}, 0.09825), respectively, with the total average number of samples required being 111851.

Using a uniform distribution, the optimum values are $N_{samples}$ = (1000, 22000, 82500, 212000) with P_{fa} = (8.3×10^{-4}, 8.3×10^{-4}, 8.3×10^{-4}, 0.09775), respectively, the

total average number of samples required, for the Gaussian distribution, being 112350. Repeating this analysis over the range of Gaussian distributions gives the contour plots in Fig. 7.2, Fig. 7.3 and Fig. 7.4.

The white shaded areas on some of the plots indicate where the uniform distribution settings exceeded the optimised settings in terms of performance. The optimisation result depends on the starting point, since the starting points were different for both systems there is a probability that the flat distribution optimised system has performance similar to the fully optimised system. Note that the flat distribution optimised system has performance only slightly greater than the fully optimised system (0.01%) in these cases.

As N_R increases, the relative error generally decreases. This is shown more clearly in Fig. 7.5 where the mean and the variance of the data in the contour plots are shown.

As can be seen, the average relative error is low. For N_R=7 the average error is less than 1.35%. The greatest error occurs for narrow distributions (distributions with a low variance) and with a low mean SNR, i.e. μ = -19 dB and σ = 1 dB, where the relative error is almost 12%.

7.3.3 CF Detector Optimised for Uniform Distribution

The greatest error in both detector architectures occurs for distributions with low variances. This would only occur in a situation where the SNR over multiple channels remains static (at the same value across all bands) over time. This situation is unlikely, in practice, though it is possible if the receiver were static, as well as the primary users. Apart from this unlikely case, the uniform distribution provides quite an effective set of operating points for a wide range of distributions.

A further advantage of the uniform distribution method, beyond the removal of the SNR

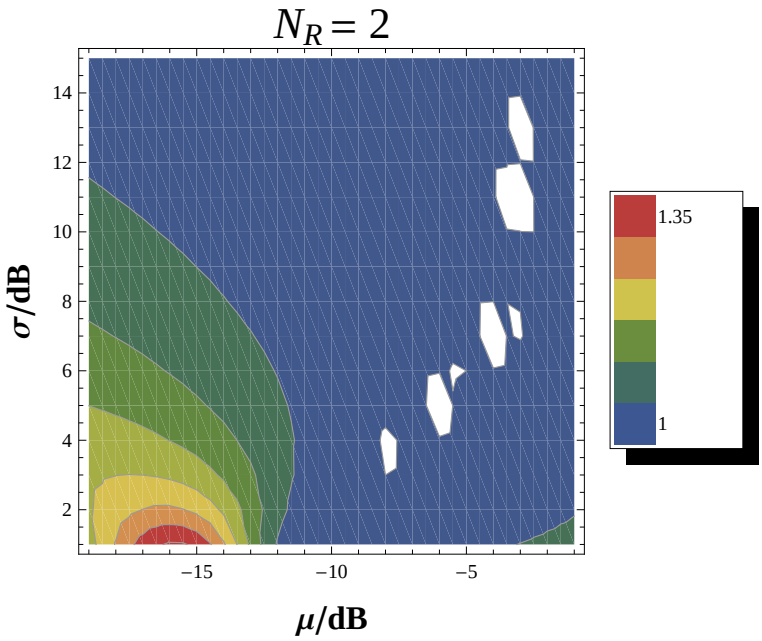

(a)

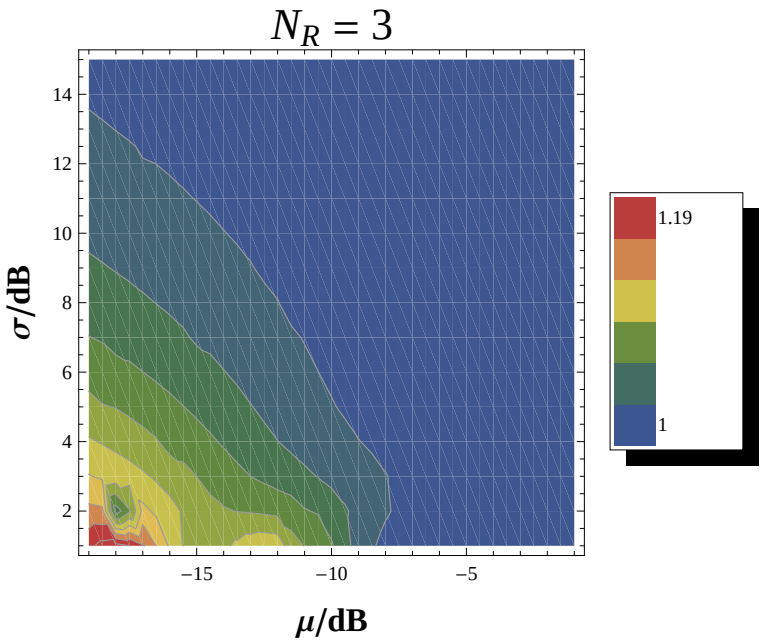

(b)

Figure 7.2: Total Number of Samples Required for Detection for CDFD Optimised for a Uniform Distribution Relative to a Directly Optimised System for (a) N_R = 2 and (b) N_R = 3

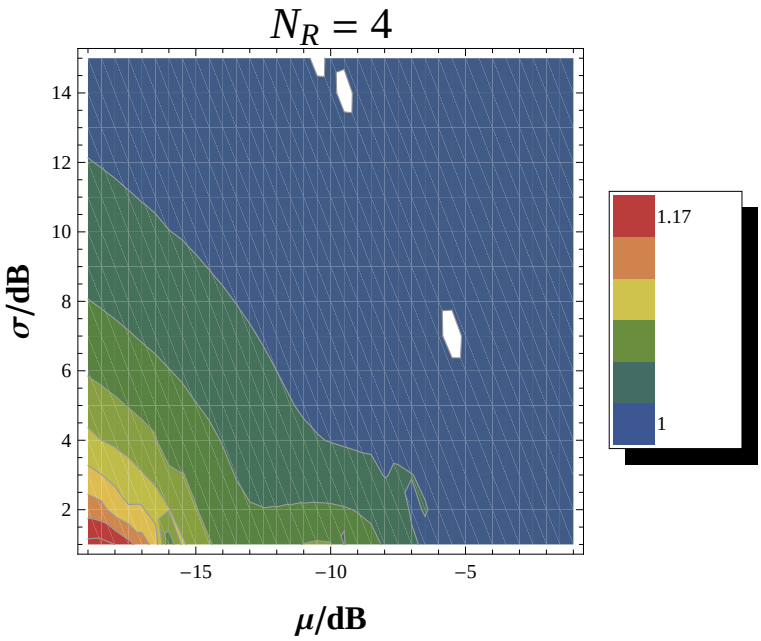

(a)

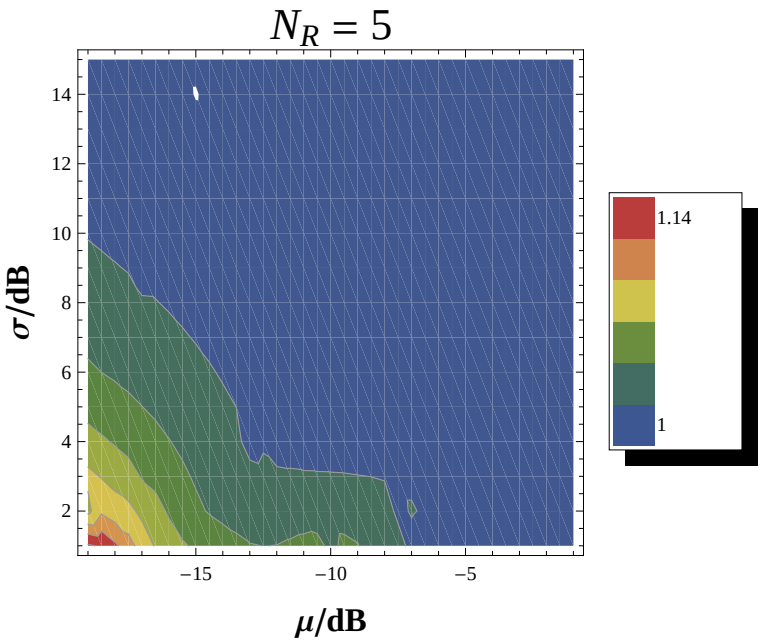

(b)

Figure 7.3: Total Number of Samples Required for Detection for CDFD Optimised for a Uniform Distribution Relative to a Directly Optimised System for (a) N_R = 4 and (b) N_R = 5

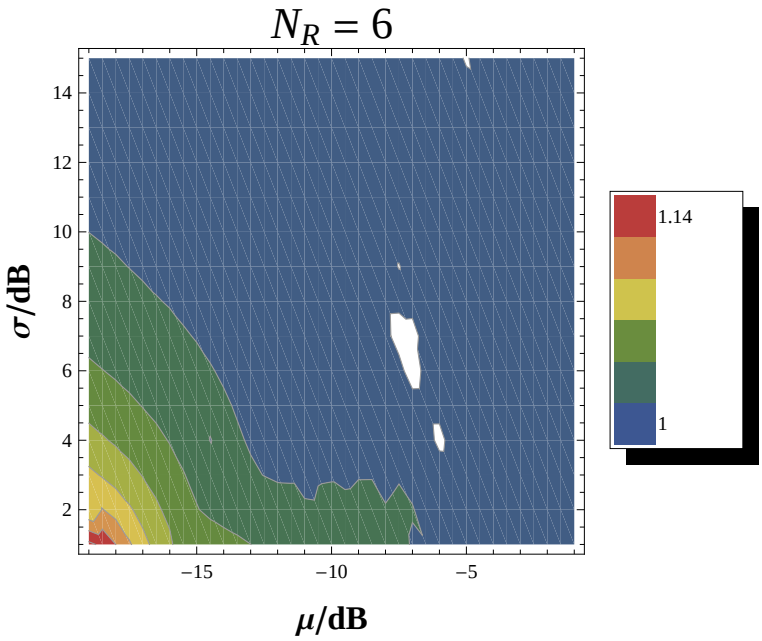

(a)

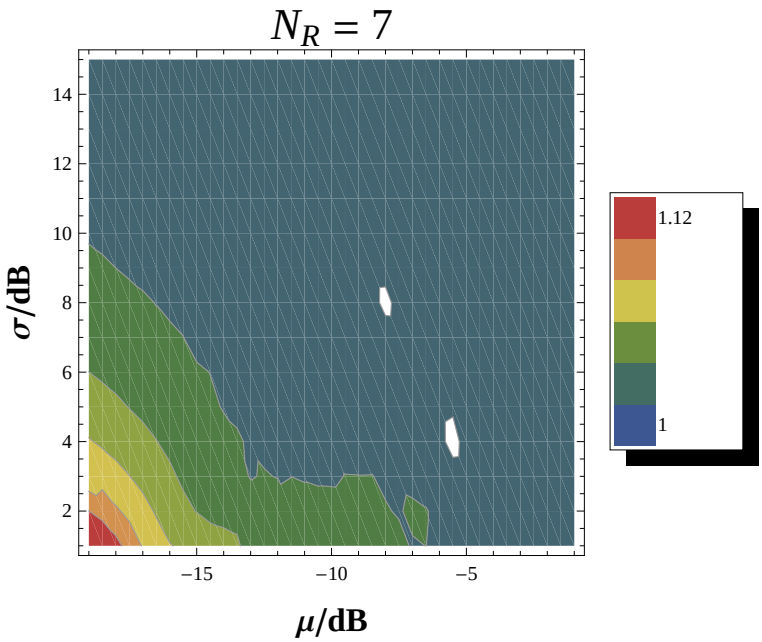

(b)

Figure 7.4: Total Number of Samples Required for Detection for CDFD Optimised for a Uniform Distribution Relative to a Directly Optimised System for (a) N_R = 6 and (b) N_R = 7

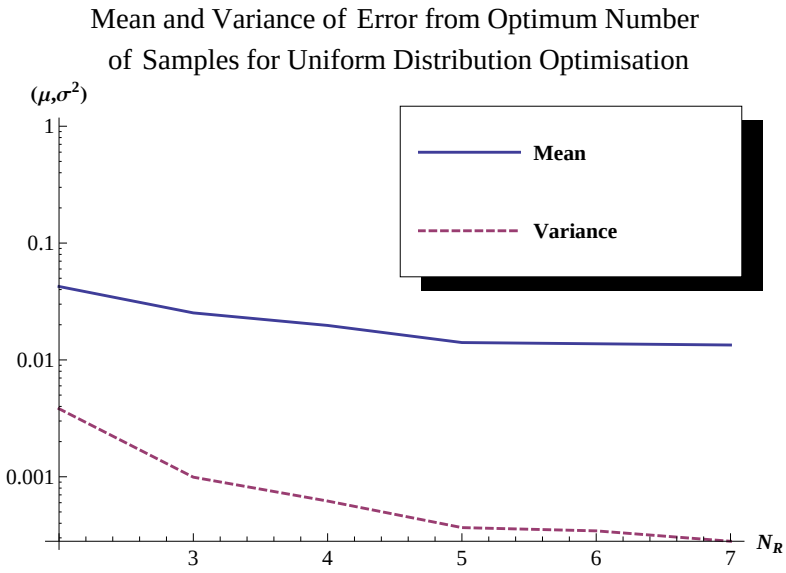

Figure 7.5: Mean and Variance of Relative Error for Uniform Distribution

distribution estimation and the associated optimisation, is the static nature of the system. Even if the underlying distribution changes dramatically, the uniform distribution method remains consistent. A continuously optimising detector would find this change in SNR detrimental to performance until the new distribution has been learned. For example, if the SNR distribution had a mean of -17dB and a standard deviation of 4dB, an optimising detector would set its parameters accordingly. If the distribution then changed in mean to -4 dB the optimum parameters would change. A CDFD optimised for a uniform distribution would require 107.8k samples, whereas optimisation with knowledge of the distribution would require 119k, for the first N_{obs} detection attempts, where N_{obs} is the number of observations, or detection attempts, when a signal is present, to meet the target $MISC$ for the detector. Therefore, the optimising detector would require approximately

10% more samples, on average, until the new distribution was learned.

Finally, the complexity of the optimisation of the CSFD is such that inline optimisation is not currently feasible (due to computational constraints); thus, the uniform distribution is the most realistic, practical option.

An interesting question that has not been answered by this work is whether there is any distribution that would cause a significant failure in the uniform distribution method? No such distribution was found in this work, though it may be possible. It is envisaged that such a distribution, if one exists, would probably be multimodal with distinct peaks. This, however, lies outside the scope of this work.

7.4 Hybrid Coarse Fine Detector

As has been noted in Chapter 5 and Chapter 6, the individual detectors, CSFD and CDFD, perform different actions on the signals received to improve the overall efficiency of the system. It is logical, then, to ask if a combination of the two architectures would perform better than either of the individual architectures. As was stated in Section 5.10, the CSFD does not discount a channel, even if the result of the coarse detection attempt is almost certainly from an occupied channel. This leads to channels that are very likely to be occupied being searched without any gain in safety for the PU. The CDFD, as stated in Section 6.9, has the opposite problem. If a result is very close to the threshold then another channel might be less likely to be occupied, i.e. no sorting is performed.

By combining the sorting action of the CSFD and the censoring action of the CDFD, the system reduces the overall time spent attempting to find a free channel, whilst not increasing the probability of interference to the PU. Here, a system with these hybrid properties is called a HCFD.

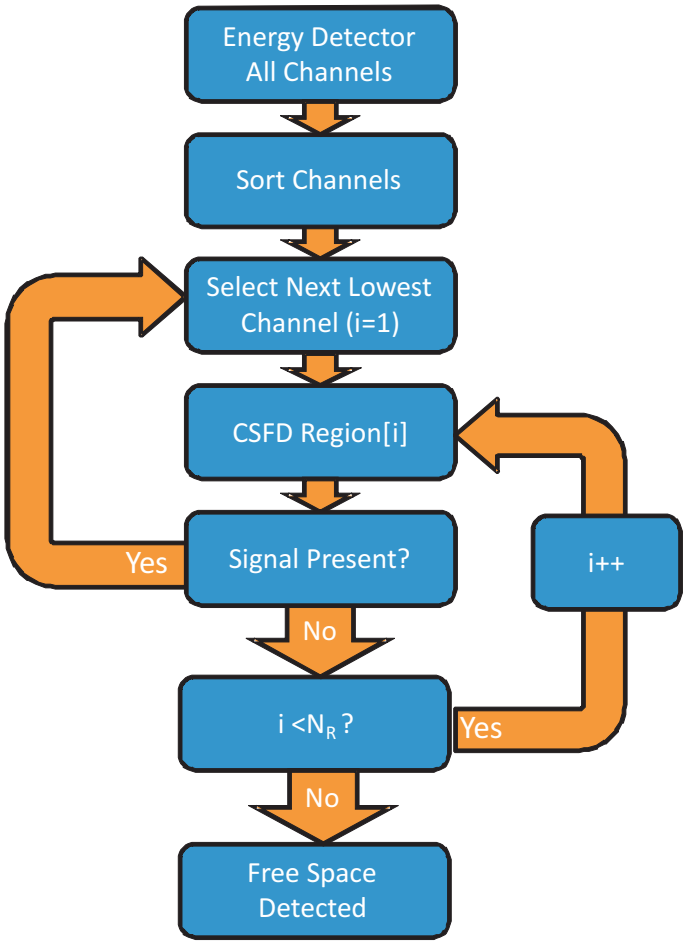

Figure 7.6: Flow Diagram for HCFD

The improvement in performance for the HCFD is not simply the multiple of both system performance gains. The CSFD and CDFD mutually counteract one another, thereby reducing the overall gain. The sorting action of the CSFD reduces the effective occupancy of the channel, which, in turn, lowers the gain from the CDFD. The CDFD reduces the cost of an incorrect choice, which lowers the gain from the CSFD. However, the overall average number of samples required is still less than either of the two other CF detectors.

The first detection attempt in the CDFD can be used for the HCFD sorting action. The

optimum number of samples for the CSFD is less than the optimum first number of samples for the CDFD. However, as the HCFD has a lower cost than the CSFD when making an incorrect decision, the optimum number of samples for sorting will be lower than the CSFD.

The HCFD could be optimised by the following method: firstly, using order statistics, an estimate of the effective occupancy for the system is found for the sorting section. The optimisation is then performed on the CDFD section of the detector using the effective occupancy previously derived. Obviously, the first detection attempt has its number of samples set by the sorting section. However, as was shown in Section 7.3, both systems perform sufficiently close to the optimum when optimised for a uniform distribution. The optimisation is more complex than either the CDFD or the CSFD optimisations. Therefore, the settings for the HCFD that are used here are the optimum ones for the CDFD.

7.5 Comparison in Simulation

To compare the three receivers with a naive detector (an energy detector without sorting), and to verify that the HCFD is indeed more efficient that the CSFD and CDFD, Monte-Carlo simulations were performed. The detectors had to find a spectrum opportunity from ten channels where the occupancy, θ, varied from 0.1 to 0.9 in steps of 0.1. The SNR was distributed as in the test distribution in Section 4.4. The results are shown in Fig. 7.7. The HCFD has the lowest average number of samples required, especially at high occupancies. At $\theta = 0.9$, the HCFD requires 355k samples compared with 390k samples for the CDFD, 627k samples for the CSFD and 1127k samples for the energy detector. The values of P_{fa} and P_{md} were also generated for each detector and the results are shown in Fig. 7.8 and Fig. 7.9, respectively.

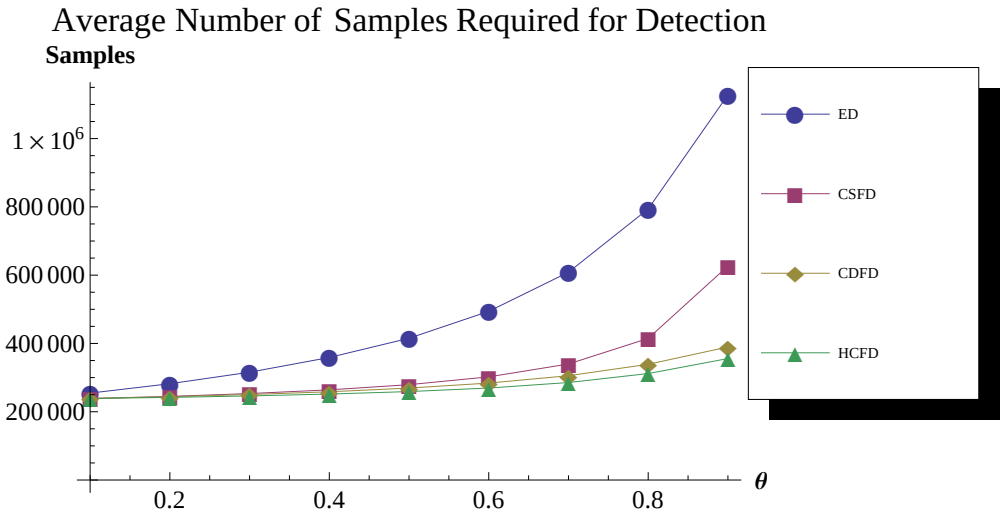

Figure 7.7: Mean Number of Samples Required for Detection for Four Detector Architectures

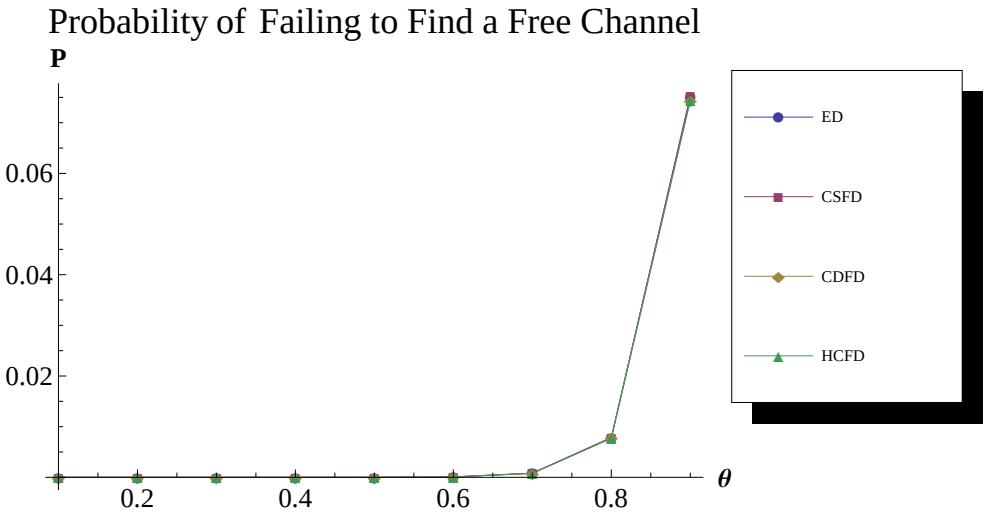

Figure 7.8: P_{fa} for the Four Detector Architectures

The four architectures have very similar false alarm and missed detection rates. The CF detectors do not cause a change in the values of P_{fa} and P_{md}, even when the average

number of samples required is significantly lower than the energy detector. Therefore, for the test distribution, it is clear that the HCFD is the best choice of all the detectors considered here. This test assumes that there is no time penalty for changing channels, i.e. that the system either has multiple PLLs or a wideband frontend. If this is not the case, then the HCFD may not be more efficient when compared to the CDFD, as the HCFD scans each channel once before beginning operation.

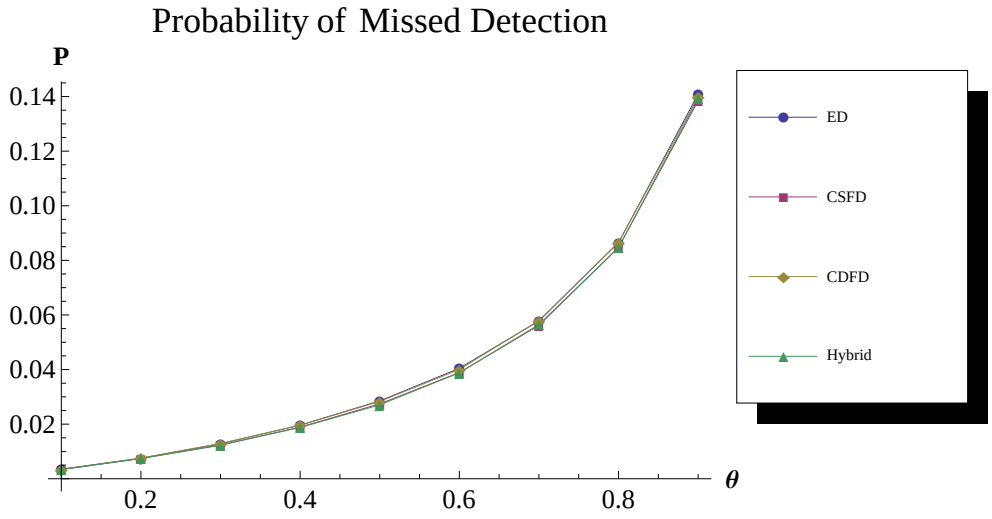

Figure 7.9: P_{md} for the Four Detector Architectures

7.6 IRIS Test

As has been stated in Section 7.2, simulation results are not a guarantee of performance in reality. To ensure that the systems are practical for implementation, the IRIS system is used to test the architectures.

7.6.1 IRIS set-up

The test was performed simultaneously on all four receiver architectures. This ensured that the test would be as fair as possible. The four architectures were implemented on one node each, with each computer and USRP frontend being identical. A picture of the IRIS receivers is shown in Fig. 7.10. Since only five nodes remained available for transmission, six channels were available as 100% occupancy was not used in the test.

Figure 7.10: Picture of IRIS Receivers at CREW Test-Bed

Each of the five nodes had a distribution of signal powers. One of the nodes had a very high average signal strength $SNR \geq 50dB$, two had a low average signal strength $10 \leq SNR \leq -10dB$ and two had a very low average signal strength $SNR \leq -15dB$, relative to the noise. A snapshot of the spectrum when all nodes were transmitting is shown in

Fig. 7.11. The frequencies used for transmission were 5010MHz to 5020MHz in steps of 2MHz and the frequency of 5008MHz was used as a free channel for noise power estimation.

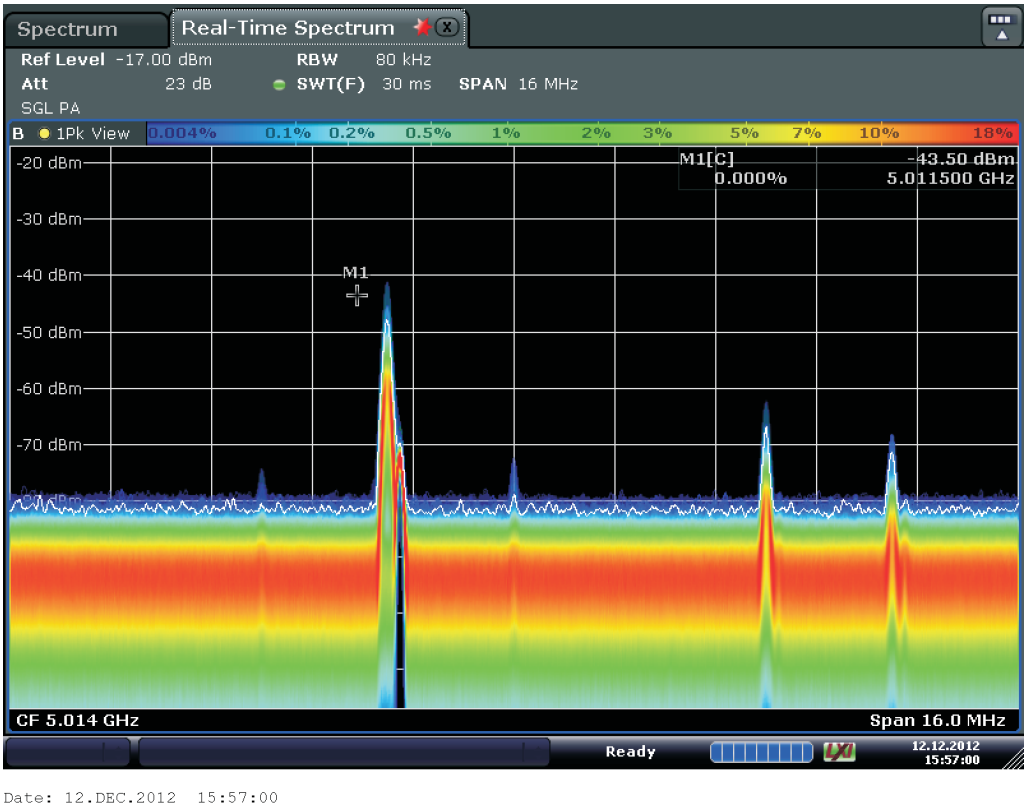

Figure 7.11: IRIS Transmitter Spectrum Snapshot

To allow the detectors to determine the occupancy of the channel and whether a false alarm or a missed detection had occurred, the occupancy of the channels was set in advance. As the detectors either sort the channels or select them randomly, this is analogous to having a random occupancy in the channels. The occupancy profiles are shown in Fig. 7.12. Finally, the SNR distributions received for each of the occupancies are shown in Fig. 7.13.

Occupancy	5.01GHz	5.012GHz	5.014GHz	5.016GHz	5.018GHz	5.02GHz
1/6	0	0	1	0	0	0
2/6	1	0	1	0	0	0
3/6	1	1	1	0	0	0
4/6	1	1	1	0	1	0
5/6	1	1	1	0	1	1

Occupied	Unoccupied
1	0

Figure 7.12: IRIS transmitter Occupancy Profile

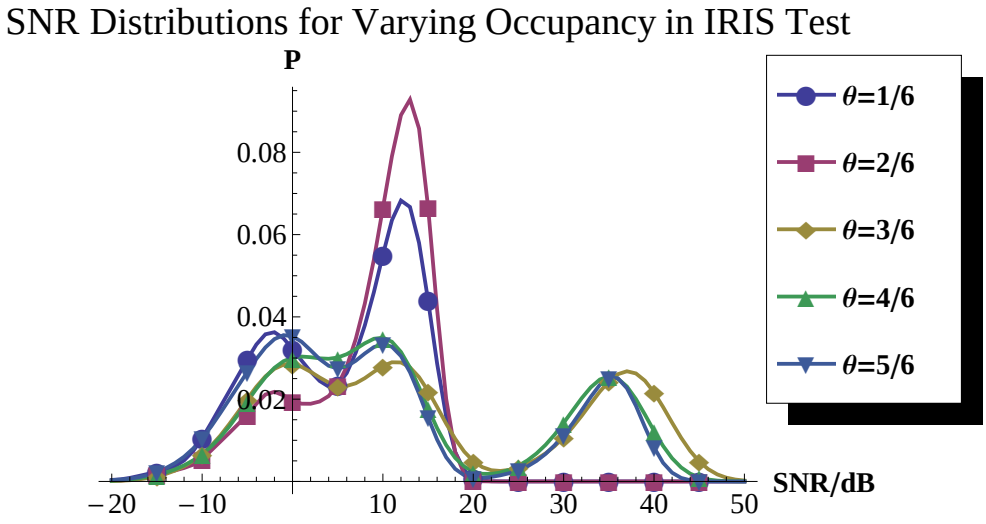

Figure 7.13: SNR Distributions

Each of the tests at different occupancies had a different SNR distribution, as the transmitters in use changed with the occupancy. There were some very strong signals, $SNR \geq 50dB$, but also some weak signals $SNR \leq -15dB$. In a real environment, it might be the case that there would be a greater number of weak signals. However, due to the lim-

itations of the USRPs as sensing devices (discussed in 7.6.2), stronger signals were used for the test.

7.6.2 IRIS results

To compare the IRIS results to simulations the environment must be also be simulated as accurately as possible. The average number of samples required from IRIS tests and simulations are shown in Fig. 7.14. The results for the average number of samples required for detection match quite well with the simulations. The HCFD detector required approximately the same number of samples as the CDFD though, for high occupancy channels ($\theta = 0.9$), the HCFD required fewer samples.

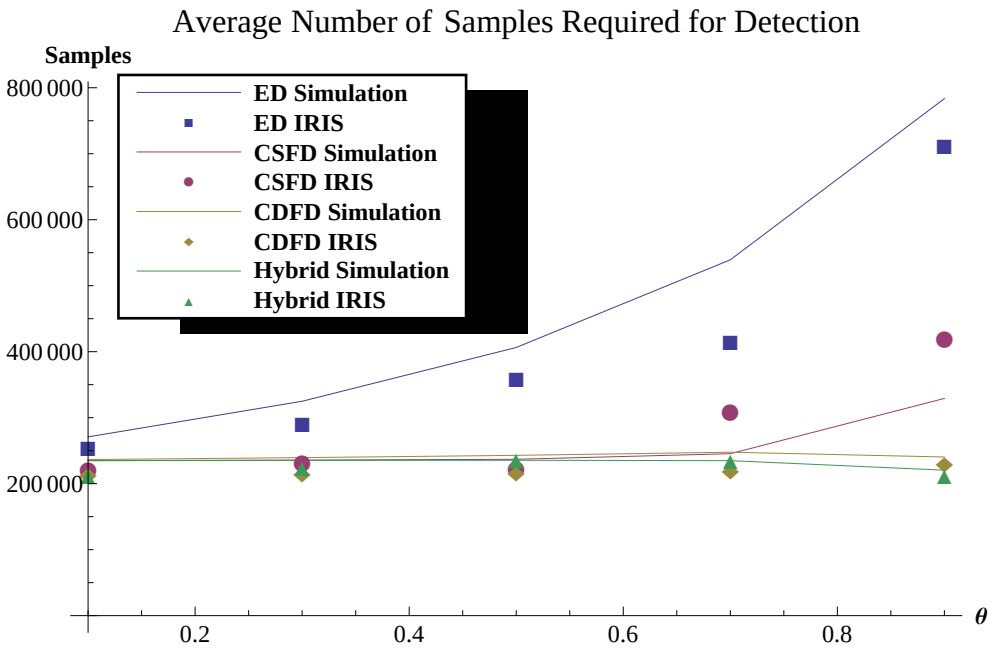

Figure 7.14: Average Number of Samples Required as a function of occupancy: IRIS vs Simulation

Shown in Fig. 7.15 are the probabilities of failing to find a free channel. This is not

exactly equal to P_{fa}, as defined here. P_{fa} is more usually defined for a single detection attempt in a single channel, this is not the case here. However, it is the real figure of merit for a detector scanning multiple bands.

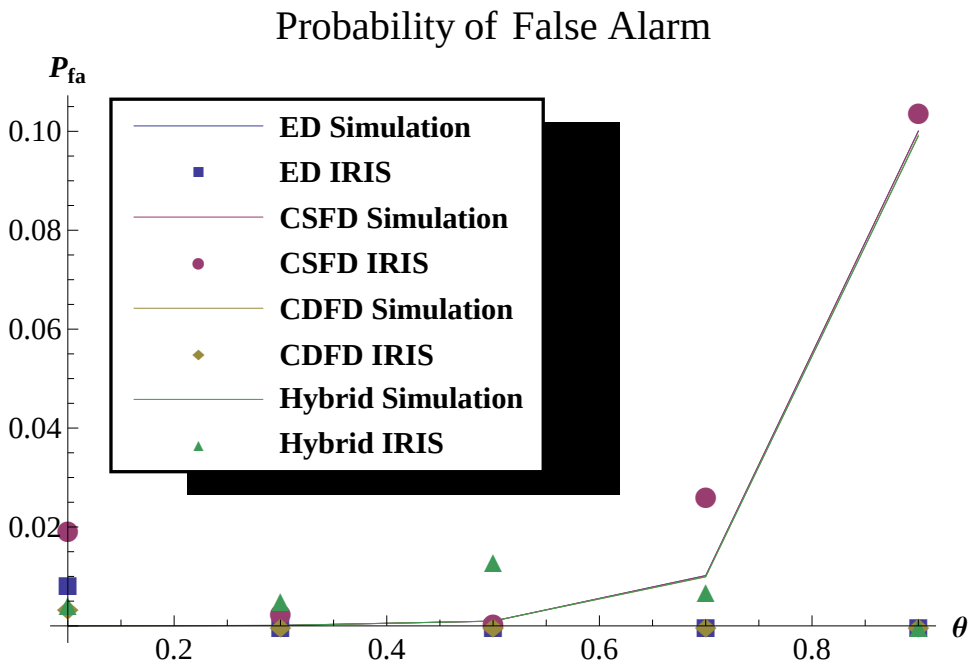

Figure 7.15: Probability of Failing to Find a Free Channel as a function of occupancy: IRIS vs Simulation

The probabilities found in the IRIS test do not match the simulations as well as the average speed gain results did. However, the probabilities of failing to find a free space are quite small. As a P_{fa} of 0.1 is deemed acceptable for the IEEE 802.22 standard, the detectors are sufficiently accurate, with only the CSFD having a probability of failing to find a free channel greater than 0.02.

Fig. 7.16 illustrates the values of P_{md} for the various architectures. The P_{md} values of some of the systems do exceed the IEEE 802.22 requirement of $P_{md} \leq 0.1$. Both the CDFD and the energy detector have $P_{md} \approx 0.2$, at $\theta = 0.8$. The other detectors do not

have $P_{md} > 0.03$ for any values of occupancy. The predicted values for the P_{md} were significantly lower ($\approx 10^{-4}$). The simulations matched the practical results well for some parameters (average number of samples required) and poorly for others (P_{md}). Some possible reasons for the mismatch are now discussed.

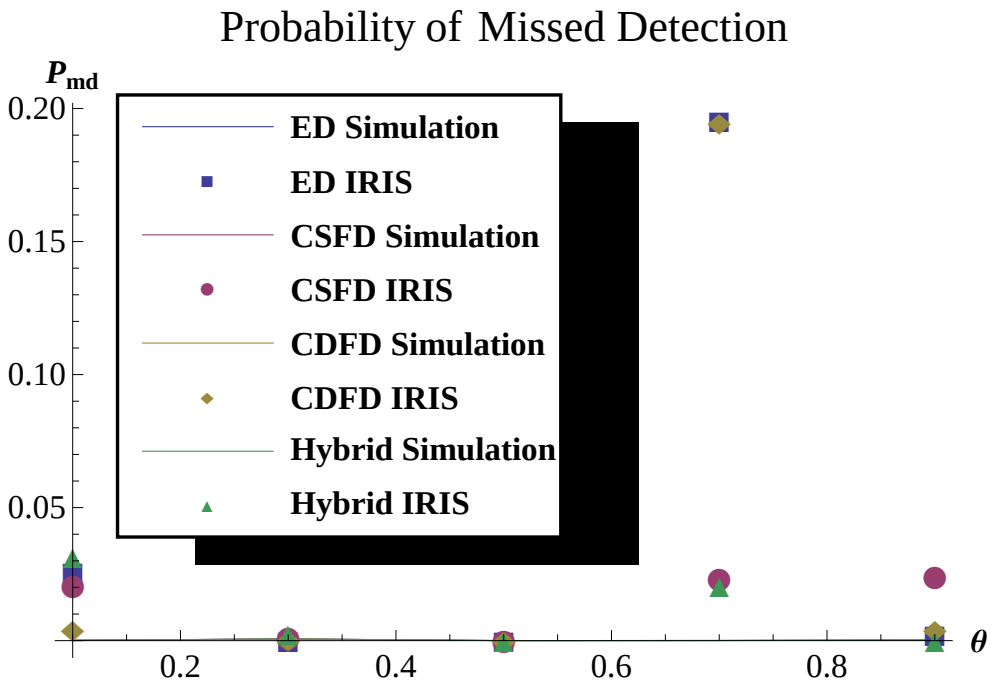

Figure 7.16: Probability of Causing Interference as a function of occupancy: IRIS vs Simulation

The CDFD requires approximately the same average number of samples for detection as the HCFD. The calibration errors, however, act to reduce this number. The P_{md}, which causes the detector to end its detection attempt early, is significantly greater in the CDFD than in the HCFD. Therefore, the average number of samples for a better calibrated CDFD would be greater.

Firstly, the number of runs that each detector had at each occupancy level was only ap-

proximately 1000. The testing required the usage of the entire CREW test-bed and, as such, could not be performed for long periods of time. If more runs could be performed, the results would most likely become more accurate and may match the simulations better, though this is not certain.

In addition, the USRPs are not designed for precise spectrum measurements. The gain of any USRP varies with time and temperature. Thus, the noise power scaling that is needed changes constantly. This was not accounted for in the simulations. In a real CR system the front-end would be more stable and the gain would vary less, allowing more accurate sensing to be performed.

In addition, the gain of the USRPs varies with frequency. An extra scaling term had to be included for each channel used by the USRP. Thus, calibration before use was required for each USRP, as the scaling was different for each node. In addition, the scaling required varied during operation, causing further calibration issues.

The primary effect of these issues is to increase the noise uncertainty in the system, though the magnitude of the corresponding noise uncertainty was not calculated or estimated. As the gain drifted with time it is impossible to say what the uncertainty is at any moment. If this was not the case then it would be possible to determine roughly the level of uncertainty as a function of the number of samples taken during the noise power estimation phase.

In a practical CR the noise uncertainty would, most likely, be lower and it is believed the results would match more closely with those predicted here by simulation.

Finally, it is noted that the HCFD had the best performance of the systems considered. It required the fewest samples, on average, for detection and did not exceed the required P_{fa} and P_{md} specifications. However, as the settings on the USRPs were so important to the operation of the system, it is likely that the HCFD radio front-end was simply better

calibrated than the other radios.

7.7 Conclusion

In this chapter, practical CR spectrum sensing was considered and the requirements for a practical CR were examined. A new architecture, the HCFD, was introduced and shown to be superior to the CSFD and CDFD in terms of average number of samples required, whilst maintaining the same P_{fa} and P_{md}.

The four architectures, CSFD, CDFD, HCFD and the energy detector, were implemented on the IRIS system. The results, for the average number of samples required, match well with the simulations. For the P_{fa} and P_{md} the results were sufficiently close, as the radio front-ends used caused difficulties for weak signal detection.

The HCFD was also shown in the IRIS tests to be a very efficient architecture. The HCFD required (at $\theta = 0.9$), on average, 212k samples when attempting to find a free channel. In contrast, the energy detector required, on average, 712k samples.

Finally, the performance of the three CF detectors, when optimised for a uniform distribution, was considered. It was shown that optimising for a uniform distribution, for the scenarios considered here, gives performance sufficiently close to the global maximum, and that this method can be used under normal operating conditions. With this result, the central question of the thesis has been answered, as stated in Section 1.2, and the work will now be concluded and the results summarised.

8

Conclusion

8.1 Summary

This thesis has investigated the use of CF spectrum sensing for DSA in CR applications. Instead of the traditional method of using a single SNR value when considering the detectors, a range of SNR values were used. This different analysis method is a powerful technique that allows a more indepth optimisation.

In Chapter 2, the current state of the art in CR was discussed. An overview of the field was

given , with special emphasis on spectrum sensing applications. Several current detector architectures for CF spectrum sensing were reviewed and discussed. Also discussed were the various test-beds currently in use and the CR system architectures used in each test-bed.

In Chapter 3, some basic theory was reviewed. The basics of energy detector operation, including the issue of noise uncertainty, were investigated. The effects of fading channels and time varying channel occupancies were also considered. Markov Chain theory was also introduced and the relevant equations governing Markov Chains shown. Finally, the IRIS system and the CREW test-bed were examined and the IRIS architecture's structure shown.

In Chapter 4, SNR PDF estimation was investigated. Various strategies for generating sample SNR PDFs were considered. In-line sensing was chosen as the most promising candidate and the advantages and disadvantages were shown. Testing and verification of the method was performed, both in simulation and on the IRIS system.

In Chapter 5, the CSFD architecture was considered. A new model of the CSFD was generated that matches Monte-Carlo simulations closely, whilst requiring significantly less (≈ 80 times) simulation time. By using order statistics to model the sorting operation and Markov Chains to model the effects of the sorting on the relevant probabilities of the fine detector, it was shown that the model predicts the CSFD performance accurately, even under fading and noise uncertainty conditions.

In Chapter 6, the CDFD architecture was considered. The characteristic equation of the CDFD was derived and three optimisation options investigated. It was shown that, by only allowing one false alarm rate to vary, performance close to the global maximum can be obtained, whilst reducing the complexity of the optimisation significantly.

In Chapter 7, the HCFD architecture was introduced. The HCFD is a combination of

both techniques, CSFD and CDFD, and has better performance than either detector. It was shown that the CSFD and the CDFD have close to optimal performance over a wide range of SNR distributions when optimised for a uniform distribution. Using this fact, the HCFD was not optimised directly, rather the parameters for the detectors optimised for the uniform distribution were chosen and the detector compared with the other architectures. This comparison was done both in simulation and using the IRIS test-bed for a practical implementation.

8.2 Contributions

In this thesis the following contributions were made:

1. *Method for generating a reliable estimate of an SNR distribution* : A method for generating a reliable estimate of an SNR distribution was developed. By using the results of the energy detector when a signal was declared present, an estimate of the SNR can be obtained, without extra processing being required. By taking a sufficient number of estimates of the individual SNR values, a reliable estimate of the distribution can be generated. A technique for measuring the accuracy of this method was proposed, the $MISE$, which compares the sequential distribution estimates to see if the distribution has settled.

   Although this method does have a bias due to false alarms reporting low SNR values, the fact that it is effectively free makes it a valuable technique. A CR using energy detection techniques incurs only a small computational cost when generating the SNR estimate. By using preloaded kernel distributions the computational cost of the distribution estimation can be greatly reduced. The system does have

difficulties in dealing with high noise uncertainty environments and there is a bias introduced by the false alarm and missed detection rates of the detector. However, even in these cases, the estimate remains sufficiently close to the actual distribution that overall system performance does not suffer greatly.

2. *CSFD model*: A model was developed of the CSFD that allowed significantly quicker simulations, compared to Monte-Carlo simulations. Using order statistics and Markov Chains to model the process, the model was approximately two orders of magnitude quicker when generating results, whilst still being very accurate, even in the presence of fading and noise uncertainty.

   This model allows the user to test the CSFD over a wider range of conditions that would be possible using Monte-Carlo simulations. Indeed, the graph shown in Fig. 7.1 could not have been generated using Monte-Carlo simulations without significant computational resources. Using the PDF method, the simulations were completed in less than two weeks on the boole cluster. Without this model the simulations would have taken 3-4 years. This method is still too slow to be used for inline optimisation however, as the CR does not have a large amount of processing power or processing time available. For design work during the development of the CR system this method can reduce the simulation overhead dramatically.

3. *CDFD characteristic equation and optimisation*: The characteristic equation of the CDFD was derived. This closed form expression predicts the performance of the CDFD accurately, even when the CDFD reuses samples. Using this equation, an optimisation scheme was proposed. This optimisation scheme was significantly less complicated than an exact optimisation, whilst still producing results close to the global optimum (within 1%) for the CDFD.

The combined equation and reduced complexity optimisation scheme can be used for inline optimisation. The computational complexity is relatively low and the entire process can be completed between sensing periods. The optimisation is only performed when the SNR distribution has changed. The CDFD does have an issue with noise uncertainty, however several possible solutions to the problem have been presented, including co-operative sensing and reducing the false alarm requirements.

4. *Uniform distribution optimisation proposed*: Instead of individual optimisation for each distribution, it was shown that assuming a uniform distribution for the SNR results in a set of parameters that are sufficiently close to the global optimum, for a wide range of SNR distributions (within 12% for the CSFD and CDFD). This removes the requirement for the CR to know the SNR distribution and for in-line optimisation, thereby reducing the complexity of the implementation.

   For the CSFD there is currently no-inline option available, therefore, the uniform distribution assumption is a good option. For the CDFD, the option of the in-line optimisation scheme allows the end-user a choice, the system can use the in-line scheme if the greatest level of performance is required. If a lower complexity or lower powered implementation is required, then the uniform distribution assumption can be used.

5. *HCFD Architecture developed*: A new architecture was developed, the HCFD, which has performance superior to that of the CSFD and CDFD, under the operating conditions considered in this thesis. At high occupancies ($\theta = 0.9$) the HCFD required approximately 57% of the samples that the CSFD required and 91% of the samples required by the CDFD, without any optimisation being performed on the

HCFD.

The exact magnitude of the improvement will depend on the operating conditions, though for all the cases considered here, the HCFD had superior performance. The HCFD does not require any extra hardware and only a slight increase in computational complexity compared to the CSFD and CDFD detectors. Therefore, it is recommended that the HCFD be used in advance of the other detector architectures studied here.

6. *Practical tests performed on architectures*: Finally, the detectors were implemented using the IRIS system on the CREW test-bed. This provided practical results for the CSFD, CDFD and HCFD that proved that the simulations are matched well by reality, within the constraints of the equipment used. For real end-user implementations, custom recievers will be designed that will provide superior performance and reduce the innacuracies that were seen in the results here.

8.3 Future Work

During the course of this work, several questions appeared that were not answered as part of this thesis. Some of these questions would make valuable additions to the work here. There was no optimisation performed on the HCFD. It was assumed that the parameters used for the CDFD would be sufficiently close to the optimum. A full test could be performed to check the optimisation. The sorting operation could be modified, as in Chapter 5, to provide an estimate of the new effective occupancy of the channel. Then the characteristic equation of the CDFD (derived in Chapter 6) could be used, with the constraint of the number of samples used in the sorting operation, to find the optimum

CDFD settings.

The SNR estimation method used was found to have an inherent bias. Some methods of reducing or removing this bias from the resulting distribution were considered but not tested. If this bias could be removed from the estimates, then the estimation technique should become significantly more reliable.

A very important question that arises from the investigation of detection attempts across multiple bands is the question of the false alarm and missed detection rates. Simply specifying P_{fa} and P_{md} targets that must be met at a certain SNR does not guarantee safety. In this work it has been noted than the P_{fa} is not the same as the probability of not finding a free channel for a detector scanning multiple bands. Similarly, the probability that a radio scanning multiple bands will interfere with one of those channels is not the same as the P_{md} of the detector. This is an important factor that should be considered for CR detectors.

Every analysis performed in this thesis assumes a single radio performing the sensing. If multiple radios are co-operating then the situation changes dramatically. A large number of co-operating radios do not require CF sensing as the diversity gain is likely to provide sufficient performance. However, a small number of co-operating radios (e.g. ≤ 5) may be able to profit from a CF sensing scheme.

Such a sensing scheme could use the architectures discussed here to reduce overall sensing time. However, it would probably be more advantageous to design the scheme directly with co-operative sensing in mind. The availability of a control channel and the allowable levels of communication between the nodes would set constraints on the system. For example, if communications between the nodes is to be minimised, then a central node can be used for the sorting detection and then broadcasts the order to the other channels. If communications are unconstrained all the nodes can perform the coarse sorting detections.

The second option would result in greater performance, especially in the presence of fading and noise uncertainty.

Even if CF sensing is not used for co-operative scheme, SNR distributions can be used to optimise the co-operative system and make it more efficient. The SNR distribution can be used to estimate the performance of a system under real operating conditions. In addition, the SNR distribution estimation technique has not been studied for co-operative sensing. The offset due to noise uncertainty and the probabilities of false alarm and missed detection will be reduced due to co-operative sensing, though the improvement has not been quantified.

The CDFD optimisation scheme implemented here uses the pattern search algorithm. The efficiency of the pattern search algorithm depends on the starting point and the initial step size. The more appropriate the selection the more efficient the search. This work does not optimise this selection of the parameters. As the optimisation of the parameters occurs when the distribution changes, the initial point should be the optimum point for the previous distribution. The initial step size should be a function of the difference between the distributions, i.e. the $MISC$. This needs to be investigated further and the exact relationship between the $MISC$ and the optimum step size found.

Finally, one major issue that was encountered during this work was the definition of the specifications for spectrum sensing. Spectrum sensing for cognitive radio applications has the requirement that the detector has a P_{fa} of less than 10% and P_{md} of less than 10% at an SNR of -21dB. However, this is not stated for a single channel or multiple channels. As cognitive radios are scanning over multiple candidate channels trying to find a free space, traditional single channel metrics are less appropriate. If it is a single channel metric, as most publications assume, then a 10% missed detection rate per channel for 10 occupied channels, results in a 65% probability of interfering with one of the channels. Instead,

new multiple scan metrics should be generated to determine the permissible interference rates for CRs operating in licensed bands.

8.4 Final Summary

This thesis has shown how a CR CF sensing scheme may be optimised for the presence of a wide range of SNRs. The detectors analysed have been shown to be significantly quicker than the naive detector and the method of using a distribution of SNR values, rather than a single value, can provide significant performance gains. A CF system optimised for a uniform distribution has performance sufficiently close to the global optimum for most practical purposes. Therefore, the uniform distribution assumption is recommended for use with the methods in this thesis to optimise any CF system scheme based on sorting channels, deciding on channels, or a combination of both, for any SNR distribution of the types considered here.

Bibliography

[1] J. Mitola, “Cognitive radio: Model-based competence for software radios,” Ph.D. dissertation, Dept. of Teleinformatics, KTH, 1999.

[2] D. Cabric, S. Mishra, and R. Brodersen, “Implementation issues in spectrum sensing for cognitive radios,” in *Signals, Systems and Computers, 2004. Conference Record of the Thirty-Eighth Asilomar Conference on*, 2004, pp. 772– 776.

[3] S. J. Shellhammer, “Spectrum sensing in IEEE 802.22,” in *IAPR Wksp. Cognitive Info. Processing*, June 2008, pp. 1–6.

[4] “From the FCC lab: Report on trends in wireless devices,” FCC, Tech. Rep., January 2011.

[5] “Cisco visual networking index: Global mobile data traffic forecast update, 2011 to 2016,” CISCO Systems, Tech. Rep., February 2012.

[6] FCC, “Spectrum policy task force report,” ET Docket No. 02-155, Tech. Rep., November 2002.

[7] D. Cabric, S. M. Mishra, D. Willkomm, R. Brodersen, and A. Wolisz, "A cognitive radio approach for usage of virtual unlicensed spectrum," in *In Proc. of 14th IST Mobile Wireless Communications Summit*, 2005.

[8] G. Staple and K. Werbach, "The end of spectrum scarcity [spectrum allocation and utilization]," *Spectrum, IEEE*, vol. 41, no. 3, pp. 48–52, March 2004.

[9] J. Mitola, "Cognitive radio: Making software radios more personal," *Personal Communications, IEEE*, vol. 6, no. 4, pp. 13–18, August 1999.

[10] 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.

[11] A. B. MacKenzie, J. H. Reed, P. Athanas, C. W. Bostian, R. M. Buehrer, L. A. DaSilva, S. W. Ellingson, Y. T. Hou, M. Hsiao, J.-M. Park, C. Patterson, S. Raman, and C. R. daSilva, "Cognitive radio and networking research at Virginia Tech," *Proceedings of the IEEE*, vol. 97, no. 4, pp. 660–688, April 2009.

[12] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 201–220, February 2005.

[13] B. Razavi, "Cognitive radio design challenges and techniques," *Solid-State Circuits, IEEE Journal of*, vol. 45, no. 8, pp. 1542–1553, 2010.

[14] I. . af", IEEE Standard.

[15] H. Zamat, "Practical implementation of sensing receiver in cognitive radios," Ph.D. dissertation, Kansas State University, 2009.

[16] G. Singh, "Optimization of spectrum management issues for cognitive radio," *Journal of Emerging Technologies in Web Intelligence*, vol. 3, no. 4, pp. 263–267, 2011.

[17] V. Blaschke, T. Renk, and F. Jondral, "A cognitive radio receiver supporting wideband sensing," in *IEEE International Conference on Communications Workshops, ICC Workshops*, May 2008, pp. 499 – 503.

[18] J. Park, Y. Hur, T. J. Song, K. Kim, J. Lee, K. Lim, C.-H. Lee, H. S. Kim, and J. Laskar, "Implementation issues of a wideband multi-resolution spectrum sensing (mrss) technique for cognitive radio (cr) systems," in *Cognitive Radio Oriented Wireless Networks and Communications. 1st International Conference on*, 2006, pp. 1–5.

[19] W. Krenik and A. Batra, "Cognitive radio techniques for wide area networks," in *42nd Design Automation Conference, Proceedings of*, 2005, pp. 409–412.

[20] F. Khozeimeh and S. Haykin, "Self-organizing dynamic spectrum management for cognitive radio networks," in *Communication Networks and Services Research Conference (CNSR), Eighth Annual*, 2010, pp. 1–7.

[21] D. O. Hebb, "Organization of behavior," *Journal of Clinical Psychology*, vol. 6, no. 3, p. 335, 1950.

[22] Q. Cai, S. Chen, X. Li, N. Hu, H. He, Y.-D. Yao, and J. Mitola, "An integrated incremental self-organizing map and hierarchical neural network approach for cognitive radio learning," in *The International Joint Conference on Neural Networks (IJCNN)*, 2010, pp. 1–6.

[23] T. C. Clancy and N. Goergen, "Security in cognitive radio networks: Threats and mitigation," in *Proc. 3rd Int. Conf. Cognitive Radio Oriented Wireless Netw. Commun. (CrownCom)*, 2008, pp. 1–8.

[24] J. Shrestha, A. Sunkara, and B. Thirunavukkarasu, "Security in cognitive radio," San Jose State University, Tech. Rep., 2010.

[25] G. S. Kasbekar and S. Sarkar, "Spectrum auction framework for access allocation in cognitive radio networks," in *MobiHoc, Proceedings of the tenth ACM international symposium on Mobile ad hoc networking and computing*, 2009, pp. 1841 – 1854.

[26] T. Xuezhi, L. Yutao, and X. Guisen, "Dynamic spectrum allocation in cognitive radio: Auction and equilibrium," in *International Forum on Information Technology and Application*, 2009, pp. 554 – 558.

[27] K. J. R. Liu and B. Wang, *Cognitive Radio Networking and Security: A Game-Theoretic View*. Cambridge University Press, October 2010.

[28] Y. Yuan, P. Bahl, R. Chandra, T. Moscibroda, and Y. Wu, "Allocating dynamic time-spectrum blocks in cognitive radio networks," in *MobiHoc, Proceedings of the 8th ACM international symposium on Mobile ad hoc networking and computing*, 2007.

[29] H. R. Karimi, "Geolocation databases for white space devices in the UHF TV bands: Specification of maximum permitted emission levels," in *IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2011, pp. 443–454.

[30] D. Gurney, G. Buchwald, L. Ecklund, S. Kuffner, and J. Grosspietsch, "Geo-location database techniques for incumbent protection in the TV white space," in *New Frontiers in Dynamic Spectrum Access Networks. DySPAN. 3rd IEEE Symposium on*, 2008, pp. 1–9.

[31] G. MacGougan, G. Lachapelle, R. Klukas, K. Siu, L. Garin, J. Shewfelt, and G. Cox, "Degraded GPS signal measurements with a stand-alone high sensitivity receiver," in *Proceedings of the National Technical Meeting of The Institute of Navigation*, 2002, pp. 1–14.

[32] D. Horgan and C. C. Murphy, "Performance limits of cooperative energy detection in fading environments," in *CogART '11 Proceedings of the 4th International Conference on Cognitive Radio and Advanced Spectrum Management*, 2011.

[33] J. Unnikrishnan and V. Veeravalli, "Cooperative sensing for primary detection in cognitive radio," in *Selected Topics in Signal Processing, IEEE Journal of*, 2008, pp. 18 – 27.

[34] V. Aalo and R. Viswanathan, "Asymptotic performance of a distributed detection system in correlated gaussian noise," *IEEE Transactions On Signals Processing*, vol. 40, no. 1, pp. 211–213, 1992.

[35] W. Zhang, R. Mallik, and K. Ben Letaief, "Cooperative spectrum sensing optimization in cognitive radio networks," in *Communications. ICC. IEEE International Conference on*, 2008, pp. 3411 – 3415.

[36] E. Visotsky, S. Kuffner, and R. Peterson, "On collaborative detection of tv transmissions in support of dynamic spectrum sharing," in *New Frontiers in Dynamic Spec-*

trum Access Networks. DySPAN. First IEEE International Symposium on, 2005, pp. 338 – 345.

[37] M. Bin Shahid and J. Kamruzzaman, "Weighted soft decision for cooperative sensing in cognitive radio networks," in *16th IEEE International Conference on Networks, ICON*, 2008.

[38] D. Chen, J. Li, and J. Ma, "Cooperative spectrum sensing under noise uncertainty in cognitive radio," in *Wireless Communications, Networking and Mobile Computing. WiCOM. 4th International Conference on*, 2008, pp. 1–4.

[39] H. Wang, X. Su, Y. Xu, X. Chen, and J. Wang, "Snr wall and co-operative spectrum sensing in cognitive radio under noise uncertainty," *Journal of Electronics (China)*, vol. 27, pp. 611–617, 2010.

[40] H. Arslan, *Cognitive Radio, SDR and Adaptive Wireless Systems*. Springer, 2007.

[41] J. Proakis, *Digital Communications*, 5th ed. Mc Graw Hill, 2008.

[42] R. Tandra and A. Sahaip, "SNR walls for signal detection," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, pp. 4–17, 2008.

[43] J. Ma and Y. Li, "A probability-based spectrum sensing scheme for cognitive radio," in *IEEE International Conference on Communications, ICC*, 2008, pp. 1938–1883.

[44] P. Aluru, J. Rajpurohit, M. Agarwal, S. V. R. K. Rao, and G. Singh, "Improvement in total sensing time of the receiver in the cognitive radio," in *Advances in Recent Technologies in Communication and Computing, International Conference on*, 2010, pp. 437–439.

[45] A. Ghasemi and E. S. Sousa, "Optimization of spectrum sensing for opportunistic spectrum access in cognitive radio networks," in *4th IEEE Consumer Communications and Networking Conference, CCNC*, 2007, pp. 1022–1026.

[46] M. Alamgir, M. Faulkner, P. Conder, and P. Smitha, "Modified criterion of hypothesis testing for signal sensing in cognitive radio," in *Proceedings of the 4th International Conference On Cognitive Radio Oriented Wireless Networks, CROWNCOM*, 2009.

[47] Y. Hur, J. Park, W. Woo, J. Lee, K. Lim, C.-H. Lee, H. Kim, and J. Laskar, "A cognitive radio (CR) system employing a dual-stage spectrum sensing technique : A multi-resolution spectrum sensing (MRSS) and a temporal signature detection (TSD) technique," in *Proceedings of, GLOBECOM*, 2006, pp. 1–5.

[48] W. Yue and B. Zheng, "A two-stage spectrum sensing technique in cognitive radio systems based on combining energy detection and one-order cyclostationary feature detection," in *WISA*, 2009, pp. 327–330.

[49] L. Luo and S. Roy, "A two-stage sensing technique for dynamic spectrum access," in *Communications,ICC. IEEE International Conference on*, 2008, pp. 1–9.

[50] L. Luo, N. M. Neihart, S. Royand, and D. J. Allstot, "A two-stage sensing technique for dynamic spectrum access," *IEEE Transactions on Wireless Communications*, vol. 8, no. 6, pp. 3028–3037, 2009.

[51] S. Maleki, A. Pandharipande, and G. Leus, "Two-stage spectrum sensing for cognitive radios," in *ICASSP2010*, 2010, pp. 2946 – 2949.

[52] H. Zamat and B. Natarajan, "Use of dedicated broadband sensing receiver in cognitive radio," in *Communications Workshops, ICC Workshops. IEEE International Conference on*, 2008, pp. 508 – 512.

[53] C. G. Ling Luo and S. Roy, "Joint optimization of spectrum sensing for cognitive radio networks," in *Global Telecommunications Conference, GLOBECOM*, 2008, pp. 1–5.

[54] H. Urkowitz, "Energy detection of unknown deterministic signals," *Proceedings of the IEEE*, vol. 55, no. 4, pp. 523–531, April 1967.

[55] M. Alouini, F. F. Digham, and M. K. Simon, "On the energy detection of unknown signals over fading channels," in *Proceedings of IEEE International Conference on Communications*, vol. 5, 2003, pp. 3575–3579.

[56] B. Vujicic, N. Cackov, S. Vujicic, and L. Trajkovic, "Modeling and characterization of traffic in public safety wireless networks," in *In Proc. of International Symposium on Performance Evaluation of Computer and Telecommunication Systems, SPECTS*, 2005, pp. 1–10.

[57] P. Pawelczak, H. Fokke, R. V. Prasad, and R. Hekmat. Dynamic spectrum access: An emergency network case study. [Online]. Available: http://www.wmc.ewi.tudelft.nl/~przemyslawp/files/aaf_dsa.pdf

[58] W. A. Gardner, "The spectral correlation theory of cyclostationary time-series," *Signal Processing*, vol. 11, pp. 13–36, 1986.

[59] 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. No.1, pp. 13–24, Jan 2008.

[60] ——, "Cyclostationary signature detection in multipath Rayleigh fading environments," in *CrownCom*. CTVR, 2007, pp. 408 – 413.

[61] ——, "Cyclostationary signatures for rendezvous in OFDM-based dynamic spectrum access networks," in *DySpan*, vol. 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, 2007, pp. 220–231.

[62] A. J. Fehske, J. Gaeddert, and J. Reed, "A new approach to signal classification using spectral correlation and neural networks," in *New Frontiers in Dynamic Spectrum Access Networks, First IEEE International Symposium on*, 2005, pp. 144–150.

[63] K. Po and J. ichi Takada, "Signal detection method based on cyclostationarity for cognitive radio," The Institute of Electronics, Information and Communications Engineers, Tech. Rep., 2007.

[64] M. Oner and F. Jondral, "Air interface recognition for a software radio system exploiting cyclostationarity," in *Personal, Indoor and Mobile Radio Communications, PIMRC . 15th IEEE International Symposium on*, 2004, pp. 1947–1951.

[65] Z. Lu, Y. Ma, and R. Tafazolli, "A first-order cyclostationarity based energy detection approach for non-cooperative spectrum sensing," in *Personal Indoor and Mobile Radio Communications (PIMRC), IEEE 21st International Symposium on*, 2010, pp. 554 – 559.

[66] V. Turunen, M. Kosunen, A. Huttunen, S. Kallioinen, P. Ikonen, A. Parssinen, and J. Ryynanen, "Implementation of cyclostationary feature detector for cognitive radios," in *Proceedings of the 4 International Conference on CROWNCOM*, 2009, pp. 1–4.

[67] S. Haykin, D. Thomson, and J. Reed, "Spectrum sensing for cognitive radio," *Proceedings of the IEEE*, vol. 97, no. 5, pp. 849–877, 2009.

[68] K. Kim, I. Akbar, K. Bae, J. sun Urn, C. Spooner, and J. Reed, "Cyclostationary approaches to signal detection and classification in cognitive radio," in *New Frontiers in Dynamic Spectrum Access Networks. DySPAN . 2nd IEEE International Symposium on*, 2007.

[69] A. Tani and R. Fantacci, "A low-complexity cyclostationary-based spectrum sensing for UWB and WiMAX coexistence with noise uncertainty," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 6, pp. 2940–2950, July 2010.

[70] S. Haykin, "The multitaper method for accurate spectrum sensing in cognitive radio environments," in *Signals, Systems and Computers, ACSSC . Conference Record of the Forty-First Asilomar Conference on*, 2007, pp. 436 – 439.

[71] B. Farhang-Boroujeny, "Filter bank spectrum sensing for cognitive radio," *IEEE Transactions On Signal Processing*, vol. 56, no. 5, pp. 1801–1811, 2008.

[72] M. A. L. Hernandez, M. A. Rojas, and P. Stoica, "New spectral estimation based on filterbank for spectrum sensing," in *Acoustics, Speech and Signal Processing, ICASSP . IEEE International Conference on*, 2008, pp. 3509 – 3512.

[73] I. Bajaj, Y. Gong, and X. J. Zhang, "Proactive spectrum sensing with probing power control in cognitive radio," in *Signal Processing and Communication Systems (ICSPCS), 4th International Conference on*, 2010, pp. 1–5.

[74] N. M. Neihart, S. Roy, and D. J. Allstot, "A parallel, multi-resolution sensing technique for multiple antenna cognitive radios," in *International Symposium on Circuits and Systems, ISCAS*, 2007.

[75] Y. Zeng and Y.-C. Liang, "Spectrum-sensing algorithms for cognitive radio based on statistical covariances," *Vehicular Technology, IEEE Transactions on*, vol. 58, no. 4, pp. 1804–1815, 2009.

[76] Y. Wang, A. Pandharipande, Y. L. Polo, and G. Leus, "Distributed compressive wide-band spectrum sensing," in *Information Theory and Applications Workshop*, 2009, pp. 178 – 183.

[77] Z. Tian and G. Giannakis, "Compressed sensing for wideband cognitive radios," in *Acoustics, Speech and Signal Processing, ICASSP . IEEE International Conference on*, 2007, pp. 1520–1549.

[78] S. Hong, "Multi-resolution bayesian compressive sensing for cognitive radio primary user detection," in *Global Telecommunications Conference, GLOBECOM, IEEE*, 2010, pp. 1–6.

[79] L. Luo and S. Roy, "Analysis of search schemes in cognitive radio," in *4th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks, SECON*, 2007, pp. 647 – 654.

[80] ——, "Modeling and analysis of detection time trade-offs for channel searching in cognitive radio networks," *IET communications*, vol. 6, pp. 819–827, 2012.

[81] K. Smitha and A. Vinod, "A multi-resolution fast filter bank for spectrum sensing," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 20, no. 7, pp. 1323–1327, 2012.

[82] W. Ejaz, N. Hasan, and H. S. Kim., "iDetection: Intelligent primary user detection for cognitive radio networks," in *Next Generation Mobile Applications, Services and Technologies (NGMAST), 6th International Conference on*, 2012, pp. 153–157.

[83] S. Geethua and G. Narayananb, "A novel high speed two stage detector for spectrum sensing," in *2nd International Conference on Communication, Computing & Security, ICCCS*, 2012, pp. 682–689.

[84] L. Z. Yi Zhang and C. Tang, "Joint detection of cyclostationarity and energy in cognitive radio," in *Intelligent Systems and Knowledge Engineering (ISKE), International Conference on*, 2010, pp. 182 – 186.

[85] P. D. Sutton, K. E. Nolan, and L. E. Doyle, "A reconfigurable platform for cognitive networks," in *CrownCom*, 2006, pp. 1–4.

[86] P. D. Sutton, J. Lotze, H. Lahlou, S. A. Fahmy, K. E. Nolan, O. Baris, T. W. Rondeau, J. Noguera, and L. E. Doyle, "Iris: an architecture for cognitive radio networking testbeds," *Communications Magazine, IEEE*, vol. 48, no. 9, pp. 114–122, 2010.

[87] E. Research. Usrp documentation. Accessed on 12-11-2013. [Online]. Available: http://files.ettus.com/uhd_docs/manual/html/

[88] C. Sokolowski, M. Petrova, A. de Baynast, and P. Mahonen, "Cognitive radio testbed: Exploiting limited feedback in tomorrows wireless communication networks," in *Communications Workshops. ICC Workshops. IEEE International Conference on*, 2008, pp. 493 – 498.

[89] S. M. Mishra, D. Cabric, C. Chang, D. Willkomm, B. van Schewick, A. Wolisz, and R. W. Brodersen, "A real time cognitive radio testbed for physical and network level experiments," in *New Frontiers in Dynamic Spectrum Access Networks, DySPAN. First IEEE International Symposium on*, 2005.

[90] D. Cabric, "Addressing the feasability of cognitive radios," *IEEE Signal Processing Magazine*, vol. 25, pp. 85–93, November 2008.

[91] D. Cabric, A. Tkachenko, and R. Brodersen, "Spectrum sensing measurements of pilot, energy, and collaborative detection," in *IEEE Military Communications Conference, MILCOM*, 2006, pp. 1–7.

[92] ——, "Experimental study of spectrum sensing based on energy detection and network cooperation," in *TAPAS, Proceedings of the first international workshop on Technology and policy for accessing spectrum*, 2006.

[93] E. Matricciani, "A model of the probability distribution of the signal-to-noise ratio estimated from ber measurements," in *Vehicular Technology Conference (VTC Fall), IEEE*, 2011, pp. 1–5.

[94] A. Wiesel, J. Goldberg, and H. Messer, "Non-data-aided signal-to-noise-ratio estimation," in *Communications. ICC 2002. IEEE International Conference on*, 2002, pp. 197 – 201.

[95] D. R. Pauluzzi and N. C. Beaulieu, "A comparison of SNR estimation techniques for the AWGN channel," *IEEE Transactions on Communications*, vol. 48, no. 10, pp. 1681–1691, October 2000.

[96] A. Clarke, "On the distribution of signal to noise ratio when estimated from a power spectrum," Weapons Systems Research Laboratory, Department of Defence, Australia, Tech. Rep., 1980.

[97] E. Kreyszig, *Advanced Engineering Mathematics*. Wiley, 2006.

[98] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*, J. L. Kelley and P. R. Halmos, Eds. Van Nostrand Reinhold, 1960, Chapter 3.

[99] P. Sutton. Dyspan 2007 demo. Accessed on 12-11-2013. [Online]. Available: http://www.youtube.com/watch?v=-ie1MBKbHEg

[100] J. Goldhirsh and W. J. Vogel, "Handbook of propagation effects for vehicular and personal mobile satellite systems: Overview of experimental and modeling results," The Johns Hopkins University, Applied Physics Laboratory, The University of Texas at Austin, Electrical Engineering Research Laboratory, Tech. Rep., 1998.

[101] A. W. Bowman and A. Azzalini, *Applied Smoothing Techniques for Data Analysis*. Oxford University Press, 1997.

[102] B. Silverman, *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 1986.

[103] [Online]. Available: http://maps.techtir.com/ireland-dtt.htm

[104] G. S. N. Raju, *Antennas and Wave Propagation*, G. S. N. Raju, Ed. Pearson Education, September 2004.

[105] M. V. S. N. Prasad, "Path loss exponents deduced from VHF & UHF measurements over indian subcontinent and model comparison," *IEEE Transactions on Broadcasting*, vol. 52, no. 3, pp. 290–298, September 2006.

[106] B. Lawton and C. C. Murphy, "Coarse-fine spectrum sensing for reduced sensing time," in *4th International Conference on Signal Processing and Communication Systems (ICSPCS)*, December 13th-15th, 2010, pp. 1–7.

[107] ——, "Minimizing the coarse-fine spectrum sensing time for cognitive radios with ideal secondary detectors subject to noise uncertainty: An analytical approach," in *4th International Conference On Cognitive Radio And Advanced Spectrum Management (COGART 2011)*, October 23th-26th, 2011.

[108] ——, "Modelling coarse-fine spectrum sensing for cognitive radio (not anymore)," *Submitted to Transactions on Wireless Communications Magazine, IEEE*.

[109] SMARAD Centre of Excellence, "S-72.333 physical layer methods in wireless communication systems," Sylvain Ranvier, Helsinki University of Technology, Tech. Rep., 2004.

[110] A. C. C. N. Balakrishnan, *Order Statistics and Inference*. Academic Press, INC, 1991.

[111] H. P. Hsu, *Theory and Problems of Analog and Digital Communications*. Schaum Outline Series, 1993.

[112] A. Leon-Garcia, *Probability, statistics, and random processes for electrical engineering*, 3rd ed., M. J.Horton, Ed. Prentice Hall, 2008, iSBN 10: 0-13-147122-8.

[113] N. Balakrishnan, "Permanents, order statistics, outliers,and robustness," *Revista Matemtica Complutense*, vol. 1, p. 7107, 2007.

[114] G. Rempala and J. Wesolowski, *Symmetric Functionals on Random Matrices and Random Matchings Problems*, D. N. Arnold and A. Scheel, Eds. Springer, 2008.

[115] R. Hooke and T. A. Jeeves, "" Direct Search" solution of numerical and statistical problems," *Journal of the ACM (JACM)*, vol. 8, no. 2, pp. 212 – 229, April 1961.

[116] J. A. Nelder and R. Mead, "A simplex method for function minimization," *Computer Journal*, vol. 7, no. 4, pp. 308–313, 1965.

[117] J. Nocedal and S. J. Wright, *Numerical Optimization*, P. Glynn and S. M. Robinson, Eds. Springer, 2000.

[118] "Boole centre for research in informatics," www.boole.ucc.ie.