ABSTRACT

Title of dissertation: DYNAMIC SPECTRUM ACCESS IN COGNITIVE RADIO NETWORKS

Thomas Charles Clancy III
Doctor of Philosophy, 2006

Dissertation directed by: Professor William Arbaugh
Department of Computer Science

Since the 1930's, the Federal Communications Commission (FCC) has controlled the radio frequency energy spectrum. They license segments to particular users in particular geographic areas. A few, small, unlicensed bands were left open for anyone to use as long as they followed certain power regulations. With the recent boom in personal wireless technologies, these unlicensed bands have become crowded with everything from wireless networks to digital cordless phones.

To combat the overcrowding, the FCC has been investigating new ways to manage RF resources. The basic idea is to let people use licensed frequencies, provided they can guarantee interference perceived by the primary license holders will be minimal. With advances in software and cognitive radio, practical ways of doing this are on the horizon. In 2003 the FCC released a memorandum seeking comment on the *interference temperature model* for controlling spectrum use. Analyzing the viability of this model and developing a medium access protocol around it are the main goals of this dissertation.

A system implementing this model will measure the current interference temperature before each transmission. It can then determine what bandwidth and power it should use to achieve a desired capacity without violating an interference ceiling called the interference temperature limit.

If a system consisting of interference sources, primary licensed users, and secondary unlicensed users is modeled stochastically, we can obtain some interesting results. In particular, if impact to licensed users is defined by a fractional decrease in coverage area, and this is held constant, the capacity achieved by secondary users is directly proportional to the number of unlicensed nodes, and is actually independent of the interference and primary users' transmissions. Using the basic ideas developed in the system analysis, *Interference Temperature Multiple Access*, a physical and data-link layer implementing the interference temperature model, was formulated, analyzed, and simulated.

Overall, the interference temperature model is a viable approach for dynamic spectrum access, and ITMA is a concrete technique implementing it. With the help of cognitive radios, we can reform spectrum policy and have significantly more room to innovate, ushering in a new era of high-speed personal wireless communications.

DYNAMIC SPECTRUM ACCESS IN COGNITIVE RADIO NETWORKS

by

Thomas Charles Clancy III

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2006

Advisory Commmittee:

Professor William Arbaugh, Chair
Professor Jonathan Katz
Professor Raymond Miller
Professor Jeffrey Reed
Professor Mark Shayman

ACKNOWLEDGMENTS

I would like to thank Bill Arbaugh, my advisor, for all his support and advice since I entered the PhD program at Maryland. Also, I'd like to thank Professor Jeff Reed, and the Virgina Tech Mobile and Portable Radio Research Group for initial research direction and invaluable conversation. Additionally, I'd like to thank Paul Kolodzy, former chair of the FCC Spectrum Policy Task Force, who originally proposed using interference temperature for dynamic spectrum management. Lastly, I would like to thank the Laboratory for Telecommunication Sciences, for financially supporting my education.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

In 1934 the US Congress created the Federal Communications Commission (FCC) to consolidate the regulation of interstate telecommunication and supercede the existing Federal Radio Commission. Among its responsibilities is the management and licensing of electromagnetic spectrum within the United States and its posessions. For example, it licenses very-high frequency (VHF) and ultra-high frequency (UHF) broadcast television (TV) stations and enforces requirements on interstation interference.

The 21st century has seen an explosion in personal wireless devices. From mobile phones to wireless local area networks (WLAN), people want to be perpetually networked no matter where they are. Services like mobile phone and global positioning system (GPS) use frequencies licensed by the FCC, while others like WLAN and Bluetooth use *unlicensed bands*.

The most popular unlicensed bands are the Industrial, Scientific, and Medical (ISM) bands at 900 MHz, 2.4 GHz, and 5.8 GHz. While setting up a home wireless network to access your broadband Internet connection does not fall within the original "ISM" definition, lack of general use of these bands prompted the FCC to loosen restrictions. Within these frequency ranges, anyone can transmit at any time, as long as their power does not exceed the band's regulatory maximum. The end

result is that the ISM bands are crowded. We now have cordless phones interfering with home audio networks interfering with the uplink from your personal didigal assistant (PDA) to your computer.

Twenty years ago, the FCC was primarily concerned with long-distance telecommunications. Managing spectrum resources typically involved guaranteeing minimal interference levels between spectrum licensees, who included radio stations, broadcast television stations, and telecommunications providers. To accomplish this, they manually survey each area and select power, frequency, and bandwidth parameters for everyone that minimize overall interference.

Mitigating interference in this new wave of personal wireless devices is a much more difficult problem. Certainly individually licensing every person's home WiFi network would help cut down on interference, but it would by no means be a scalable solution. Rather than manual, static spectrum allocation, the FCC needs an automatic, distributed, and dynamic approach to managing radio frequencies in order to enable better spectrum sharing.

Such a scheme would define at least two classes of spectrum users. The first would be primary users who already posess an FCC license to use a particular frequency. The second would be secondary users consisting of unlicensed users. Other classes could also be defined, prioritizing users' access to the frequency band. Primary users would always have full access to the spectrum when they need it. Secondary users could use the spectrum when it would not interfere with the primary user.

A motivating example is the current broadcast television frequency bands. Of

the 68 channels, on average only 8 channels are used in any given TV market, or roughly 12%. Were unlicensed devices allowed to coexist with broadcast television, we would have an additional 350 MHz of prime spectrum real estate.

To this end, many approaches to dynamic spectrum allocation have been proposed. In the mid-to-late 1990s, Satapathay and Peha at Carnegie Mellon did some of the initial foundational work on spectrum sharing [27, 28, 29]. They showed that from a queueing theory perspective, independent networks could achieve an overall better capacity by cooperating. Examining the FCC-proposed "listen before talk" approach, they showed such greedy algorithms could lead to very poor spectrum utilization, and developed a system that imposed artificial penalties on greedy algorithms to keep them constrained.

Much research was conducted as a part of the *Dynamic Radio for IP-Services in Vehicular Environments* (DRiVE) and *Spectrum Efficient Uni and Multicast Services over Dynamic multi-Radio Networks in Vehicular Environments* (OverDRiVE) projects starting in 2000. Their main concern was video content delivery to vehicles [15, 22, 33]. The OverDRiVE architecture involves partitioning spectrum in space, frequency, and time. Each Radio Access Network (RAN) would be allocated certain blocks by a central authority in response to their predicted capacity needs.

Later, the CORVUS project at Stanford used similar ideas [4, 5]. They create a channelized spectrum pool from unused licensed spectrum, and have algorithms to allocate it efficiently.

Researchers at Rutgers [14, 24] have proposed schemes that use a "spectrum server" to allocate resources. In addition to the standard allocation ideas, they've

added a market model driven by resource pricing and utility functions.

OverDRIVE and CORVUS are centralized, requiring someone to decide who should use which spectral resources at what time, while guaranteeing minimal interference to licensed devices. While achieving good results, current politics involved in frequency licensing would make adopting such an approach unlikely. The need for a central authority hampers feasible deployment.

More recently, research has begun on distributed techniques for dynamic spectrum allocation, where no central spectrum authority is required. Decentralized approaches may be less efficient, but require much less cooperation. Game theoretic aspects of this are studied in [21].

Of the decentralized approaches, some require control channel communication between devices. They make local, independent decisions about how to best communicate, and use these low-bandwidth side channels to negotiate communications paramters [16, 35]. An open question still exists though – what if your control channel is overcome by interference? How do you negotiate a transition to a new frequency?

Challapali, et al, assumes that licensed digital signals will likely have a periodic pattern in how they transmit, and spectrum access opportunities can be predicted and exploited [6].

This dissertation proposes an entirely new concept for dynamic spectrum allocation. Our radio nodes treat licensed users, other unlicensed radio networks, other unlicensed nodes within the same network, interference, and noise all as *interference* affecting the signal-to-interference ratio (SIR). Higher interference yields lower SIR,

which means lower capacity is achievable for a particular signal bandwidth. Radio nodes search for gaps in frequency and time where the measured interference is low enough to achieve communication at a target capacity, subject to an overall interference constraint. A major distinction is that all other proposed schemes *avoid* licensed signals, while we try to *coexist* with them.

Chapter 2 covers background material, including an overview of cognitive radio technology and an introduction to information theory. Chapter 3 describes the Interference Temperature Model (ITM) and derives information theoretic bounds on its performance. Chapter 4 details Interference Temperature Multiple Access (ITMA), a medium access control scheme implementing the ITM. Chapter 5 extends ITMA to support arbitrarily-shaped signal power spectrums in an effort to take better advantage of the ITM. Chapter 6 concludes and discusses future work.

Chapter 2

Foundations

This chapter covers basic background material. Its goal is to bring the average reader up to speed on some of the history and analytical techniques that will be later employed.

2.1 Cognitive Radio Evolution

Over the past 15 years, notions about radios have been evolving away from pure hardware-based radios to radios that involve a combination of hardware and software. In the early 1990s, Joseph Mitola introduced the idea of *software defined radios* (SDRs) [17]. These radios typically have a radio frequency (RF) front end with a software-controlled tuner. Baseband signals are passed into an analog-to-digital converter. The quantized baseband is then demodulated in a reconfigurable device such as a field-programmable gate array (FPGA), digital signal processor (DSP), or commodity personal computer (PC). The reconfigurability of the modulation scheme makes it a software-defined radio.

In his 2000 dissertation, Mitola took the SDR concept one step further, coining the term *cognitive radio* (CR) [19]. CRs are essentially SDRs with artificial intelligence, capable of sensing and reacting to their environment. Figure 2.1 graphically contrasts traditional radio, software radio, and cognitive radio.

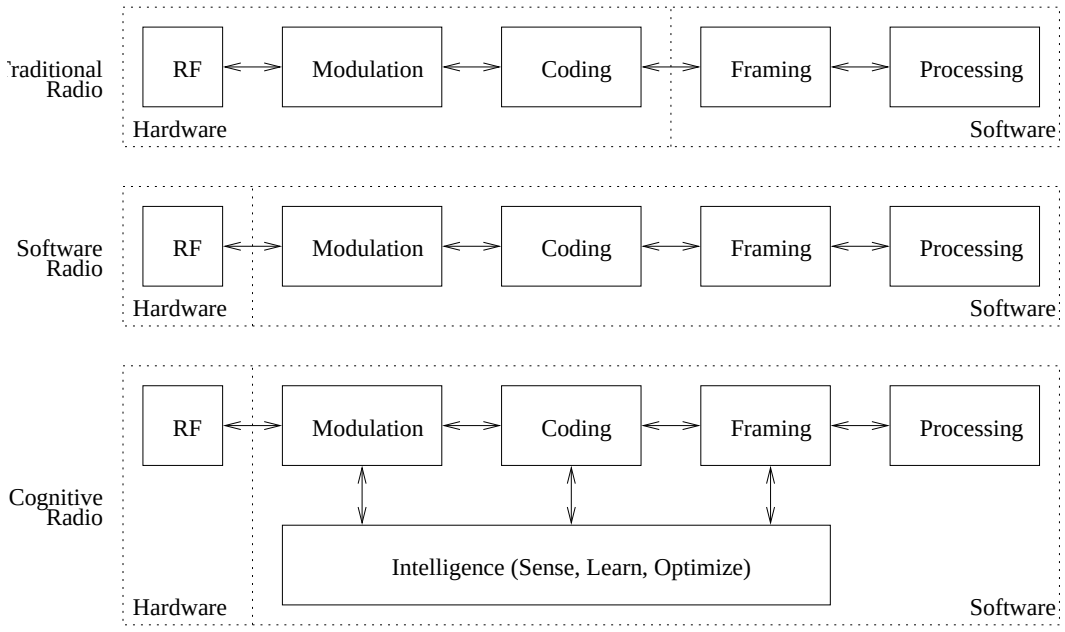

Figure 2.1: Logical diagram contrasting traditional radio, software radio, and cognitive radio

In the past few years, many different interpretations of the buzz word "cognitive radio" have been developed. Some of the more extreme defintions might be, for example, a military radio that can sense the urgency in the operator's voice, and adjust QoS guarantees proportionally. Another example is a mobile phone that could listen in on your conversations, and if you mentioned to a friend you were going to hail a cab and ride across town, it would preemptively establish the necessary cell tower handoffs [20].

Though more representative of Mitola's original research direction, these interpretations are a bit too futuristic for today's technology. A more common definition restricts the radio's cognition to more practical sensory inputs that are aligned with typical radio operation. A radio may be able to sense the current spectral environment, and have some memory of past transmitted and received packets along with

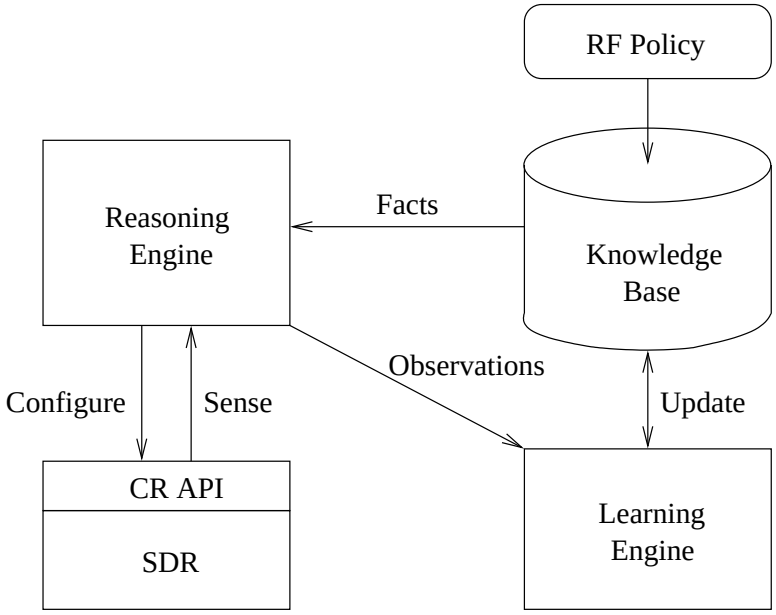

Figure 2.2: Functional portions of a cognitive radio, representing reasoning and learning capabilities

their power, bandwidth, and modulation. From all this, it can make better decisions about how to best optimize for some overall goal. Possible goals could include achieving an optimal network capacity, minimizing interference to other signals, or providing robust security or jamming protection.

Another contentious difference in interpretation has to do with drawing the line between SDR and CR. Often times, frequency agile SDRs with some level of intelligence are called CRs. However, others believe that SDRs are just a tool in a larger CR infrastructure. Remote computers can analyze SDR performance and reprogram them on the fly. For example, this remote intelligence could decide none of the SDRs' modulation schemes are sufficient for their current environment. It could create a new scheme on the fly, generate hardware description language (HDL) and new FPGA loads, and reload them over the network to add this new functionality.

Figure 2.2 shows functional components of a more concrete cognitive radio architecture. The SDR is accessed via a CR application programming interface (API) that allows the CR engine to configure the radio, and sense its environment. The policy-based reasoning engine takes facts from the knowledge base and information from the environment to form judgements about RF spectrum accessing opportunities. In addition to a simple policy-based engine, a learning engine observes the radio's behavior and resulting performance, and adjusts facts in the knowledge base used to form judgements.

A fundamental problem with a system like this is its complexity. Can the proposed learning and reasoning be done in near real time, to keep up with an ever changing RF environment? Can we come up with a simple set of metrics that can perform well without being overly computationally complex? This research aims to identify such metrics and control algorithms for implementing them. We do not overly delve into the artificial intelligence aspects of how reasoning and learning can be done in a generalized way, but apply basic principles to the metrics and algorithms we implement.

Seeing the advances in smart radio technology, the FCC began researching ways in which CRs could use licensed bands, provided they didn't interfere with existing licensees. A motion to allow operations to operate was recently approved and adopted by the FCC [11], and allows cognitive radios to operate in certain frequency bands. The FCC also proposed the *interference temperature model* [10], the instantiation of which is the subject of much of this dissertation.

2.2 Information Theory Primer

Information theory is a field that began development in the late 1940s as a result of Claude Shannon's research into communication systems. He proved fundamental limits on achievable data rates in digital communications systems. Since its inceptions, however, the basic concepts of information theory have diffused into many unrelated research fields [9, 30].

The fundamental building block of information theory is *entropy*, which measures the number of information bits conveyed by a random variable. In particular, if $\mathcal{X}$ is a discrete random variable over a space of size n, and notationally

$$\mathbf{P}[\mathcal{X} = x_i] = p_{\mathcal{X}}(x_i) \tag{2.1}$$

then the entropy $H(\cdot)$ is defined as

$$H(\mathcal{X}) = \sum_{i=1}^{n} -p_{\mathcal{X}}(x_i) \log_2 p_{\mathcal{X}}(x_i) \tag{2.2}$$

In order to maximize the entropy function, we must have $p_i = 1/n$, or a uniform probability distribution.

Next, consider a communications channel. Let $\mathcal{A}$ be a random variable representing a transmitted symbol. Let $\mathcal{B}$ represent the received signal. In most real cases, $\mathcal{A} \neq \mathcal{B}$ because noise and interference affects our communication systems.

Assuming $\mathcal{A}$ and $\mathcal{B}$ are not independent, then communication is possible. To quantify it, we define a measure called *mutual information*. Mathematically, it is

defined as

$$I(\mathcal{A};\mathcal{B}) = \sum_{i,j} p_{\mathcal{A},\mathcal{B}}(a_i, b_j) \log_2 \left(\frac{p_{\mathcal{A},\mathcal{B}}(a_i, b_j)}{p_{\mathcal{A}}(a_i) p_{\mathcal{B}}(b_j)} \right) \tag{2.3}$$

A more convenient and perhaps intuitive notation for mutual information is in terms of conditional entropy, and measures the randomness of $\mathcal{A}$ assuming we know $\mathcal{B}$ equals a particular value. It is defined in terms of conditional probabilities as

$$H(\mathcal{A}|\mathcal{B}) = \sum_{i=1}^{n} \sum_{j=1}^{n} -p_{\mathcal{A}|\mathcal{B}}(a_i|b_j) \log_2 p_{\mathcal{A}|\mathcal{B}}(a_i|b_j) \tag{2.4}$$

Using this, we can redefine mutual information as

$$\begin{aligned} I(\mathcal{A};\mathcal{B}) &= H(\mathcal{A}) - H(\mathcal{A}|\mathcal{B}) \\ &= H(\mathcal{B}) - H(\mathcal{B}|\mathcal{A}) \end{aligned} \tag{2.5}$$

In words, the mutual information between $\mathcal{A}$ and $\mathcal{B}$ is the entropy of $\mathcal{A}$ minus the entropy of $\mathcal{A}$ given full knowlege of $\mathcal{B}$. So, consider our degenerative case where $\mathcal{A}$ and $\mathcal{B}$ are independent. There $H(\mathcal{A}|\mathcal{B}) = H(\mathcal{A})$ because knowledge of $\mathcal{B}$ tells us nothing about $\mathcal{A}$. As a result,

$$\begin{aligned} I(\mathcal{A};\mathcal{B}) &= H(\mathcal{A}) - H(\mathcal{A}|\mathcal{B}) \\ &= H(\mathcal{A}) - H(\mathcal{A}) \\ &= 0 \end{aligned} \tag{2.6}$$

Next, we define the idea of *channel capacity* [30]. This represents the maximum amount of information that can be conveyed through a communications channel.

From an information theoretic perspective, a communications channel is responsible for passing data between two points, and will likely add some sort of noise to the original signal. Thus if the value of our transmitted signal is distributed as $\mathcal{A}$, the output of the channel is some function of the original signal. We specify the channel output as $\mathcal{B} = f(\mathcal{A})$ where $f(\cdot)$ specifies a probability distribution for $\mathcal{B}$ given $\mathcal{A}$.

Obviously, the more mutual information between $\mathcal{A}$ and $\mathcal{B}$, the more likely we can decode $\mathcal{A}$ from $\mathcal{B}$. If $I(\mathcal{A};\mathcal{B}) = 0$, then communication is impossible because the output of our channel is statistically independent of the channel input.

Assuming the mutual information is nonzero, the goal is to then determine exactly what data rates can be achieved by the channel. The goal is to determine the optimal probability distribution over $\mathcal{A}$ to maximize $I(\mathcal{A};\mathcal{B})$, when $\mathcal{B} = f(\mathcal{A})$. Mathematically, this can be written

$$C = \max_{p_{\mathcal{A}}(a_1)\ldots p_{\mathcal{A}}(a_n)} I(\mathcal{A};\mathcal{B}) \tag{2.7}$$

Thus, we select the optimal probability distribution for $\mathcal{A}$ to maximize the mutual information between the channel input and output.

For our purposes, we'll mostly focus on the Shannon-Hartley Law [30], which states that for an additive, white, Gaussian noise (AWGN) channel with bandwidth B, signal power S, and noise power N, the capacity C is

$$C = B\log_2(1 + S/N) \tag{2.8}$$

Assuming B is in Hertz and S and N are in Watts, C will be in bits per second. Since this theorem is the primary relation used to associate bandwidth, capacity, and power, here we formally define each.

Bandwidth is measured in cycles per second, or Hertz (Hz). It represents the difference between the upper and lower frequencies used in an RF transmission. The "edges" of a signal in the spectral domain are somewhat subjective, and often subject to the quality of a radio's bandpass filters. Thus, we will define it in terms of the baseband radio's symbol rate. If we are transmitting M-bit data symbols at R bits per second, the resulting bandwidth is M/R Hertz. Certainly spread spectrum schemes will complicate this definition, and as a result this describes the *narrowband* bandwidth. See Section 2.3 and Chapter 5 for more information on spread spectrum.

Variables S and N are our signal and noise powers, respectively. The signal power S is the average power of our signal, averaged across bandwidth B.

Noise N is caused by background electromagnetic radiation. Its value can be computed from the thermodynamic relationship

$$N = B\,k\,T_N \tag{2.9}$$

where B is our bandwidth in Hertz, $k = 1.38 \cdot 10^{-23}$ is Boltzman's constant, and T_N is the temperature of our environment in Kelvin. Noise can be artificially increased through a number of sources, the most important of which is interference, which we will discuss at length in Chapter 3.

Ideally, and in most typical scenarios, noise is AWGN. It's additive because

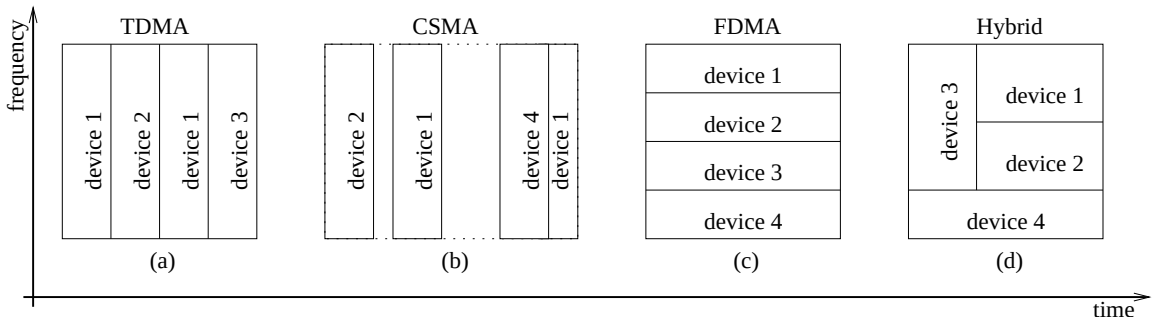

Figure 2.3: Four multiple access schemes, showing how each operates differently in the frequency and time domains.

noise is added to our signal, not multiplied. The term *white* means that the noise power at two times t_1 and t_2, where $t_1 \neq t_2$, is statistically independent. Gaussian noise simply has values distributed according to a Gaussian random process. This is generally a reasonable assumption since noise is additive and coming from many sources simultaneously. The central limit theorem says that the total noise will have a Gaussian distribution.

Later on, we will use the Shannon-Hartley theorem in the case where N is actually colored interference. It will be shown that in our environment, the bounds from the Shannon-Hartley theorem still hold.

2.3 Multiple Access

In the previous section, we were primarly concerned with two-party communications. However in most real environments, there are many users all simultaneously trying to communicate with each other. In order to support everyone, there must be a way to share the communications channel.

To achieve this, we use a multiple access scheme for multiplexing users' com-

munications. This multiplexing can be done in time, frequency, or code. Figure 2.3 graphically shows most of them, and we describe each in detail below.

In *time-division multiple access* (TDMA), users are multiplexed in the time domain, each being allocated a certain window in which to communicate. Time is typically segmented evenly into short windows, and each device in the network is assigned recurring time slots when they are scheduled to transmit. This scheduling typically requires a centralized controller in the network with knowledge of the capacity needs of each device.

A slight variant of TDMA is *carrier-sense multiple access* (CSMA). CSMA also multiplexes in time, but it does so in a less organized manner. In particular, when a device has data to transmit, it first listens. If no other devices are transmitting, it transmits. If the channel is in use, it waits until the channel is idle. A major problem in this scenario is devices transmitting simultaneously, causing interference and data loss. The advantage to CSMA is that no centralized controller is needed, but the disadvantage is that communications resources are often used inefficiently.

Moving on to the frequency domain, we have *frequency-division multiple access* (FDMA). Here, each device is assigned a frequency for communication, much like frequency modulation (FM) radio. FDMA also suffers from inefficiency, since users typically cannot switch frequencies on the fly to use idle channels if they have a lot of data to send.

There are several hybrid approaches that combine TDMA and FDMA. A centralized controller schedules data transmissions in both the frequency and time domain. A common example of this is called *orthogonal frequency division mul-*

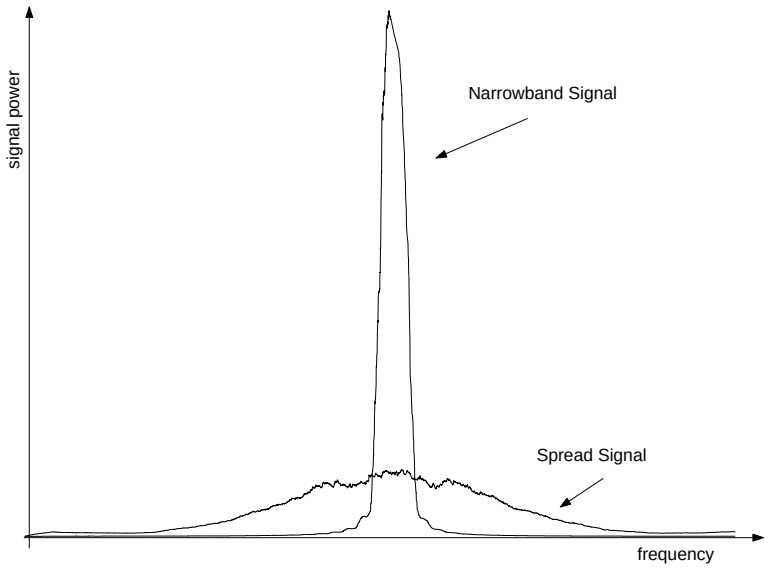

Figure 2.4: Power spectra for a narrowband signal both before and after it has been spread.

tiple access (OFDMA), which is based on *orthogonal frequency division multiplexing* (OFDM).

The final technique, called *code division multiple access* (CDMA) is more difficult to explain because it multiplexes neither in the time nor frequency domain. To understand CDMA, a basic understanding of direct sequence spread spectrum (DSSS) communications is required.

In DSSS, before being transmitted, analog waveforms are multiplied by a high-frequency pattern from the set $\{-1, 1\}^{\infty}$ called a spreading code. The effect is that the signal's power is "spread" over a larger bandwidth. We can see this illustrated in Figure 2.4. The narrowband signal is multiplied by a random stream of ± 1 at rate 15 times faster than the original signal. This creates a spread signal 15 times wider and 15 times weaker. The receiver can simply multiply the spread signal by the same spreading code to recover the original narrowband signal.

With CDMA, all devices implement DSSS and communicate simultaneously at the same frequency, however devices are all assigned unique, mutually orthogonal spreading codes. Only with the right spreading code can a receiver transform the spread signal into the original narrowband signal. The orthogonality helps minimize noise introduced into the narrowband signals. Thus at any given time, many users could be transmitting data, creating a wideband mass of spectral energy. Receivers use spreading codes to then recover the desired narrowband signal from the mass.

One implementation issue for CDMA involves receivers knowing which spreading code to use. In most systems today, all wireless devices communicate only with a centralized controller, and the controller has access to all the spreading codes. For distributed, ad hoc environments, the idea of receiver-oriented codes can be used [23], where each receiver has an assigned code, and communications are essentially addressed using particular spreading codes.

Additionally, there has been much research into techniques for despreading DSSS signals without knowledge of the original spreading code [1, 12]. Such approaches generally require short, repeating spreading codes and cyclostationary channels.

Most of these schemes will yield similar overall channel capacity, as seen by the Shannon-Hartley law. For example, in FDMA, the total bandwidth B will be subdivided into n frequency bands, yielding a fractional C/n capacity for each user. In TDMA, each transmitter will have a duty cycle of $1/n$, giving them a total capacity of C/n.

It is important to notice that we can utilize the Shannon-Hartley theorem

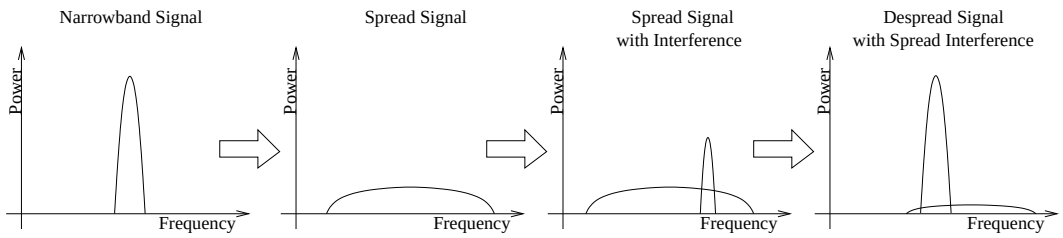

Figure 2.5: Four-step process of spreading and despreading a signal. Notice that since spreading and despreading are symmetric operations, interference power is distributed at whitened.

when measuring the capacity of CDMA in the presense of interference. Interference is inherently non-white, but the CDMA despreading process whitens interference, as long as there is no correlation between the interference and the spreading code.

The multiple access protocols and algorithms described in Chapter 4 utilize CDMA.

Overall, this chapter should have given the average reader some background in the topics discussed in this dissertation. For a more complete treatment of topics presented, see [3, 9, 18, 25, 31, 34].

Chapter 3

Interference Temperature Model

This chapter aims to analyze the interference temperature model by establishing mathematical models for the interference interactions between primary and secondary users of a particular bandwidth at a particular frequency. Using these models, probability distributions on interference are developed. From these, we can quantify both service impact on the licensee, and also achievable capacity for the underlay network. This chapter encompasses research published in [7, 8].

3.1 Interference Temperature

The concept of interference temperature is identical to that of noise temperature. It is a measure of the power and bandwidth occupied by interference. Interference temperature T_I is specified in Kelvin and is defined as

$$T_I(f_c, B) = \frac{P_I(f_c, B)}{kB} \tag{3.1}$$

where $P_I(f_c, B)$ is the average interference power in Watts centered at f_c, covering bandwidth B measured in Hertz. Boltzmann's constant k is $1.38 \cdot 10^{-23}$ Joules per Kelvin degree.

The idea is that by taking a single measurement, a cognitive radio can com-

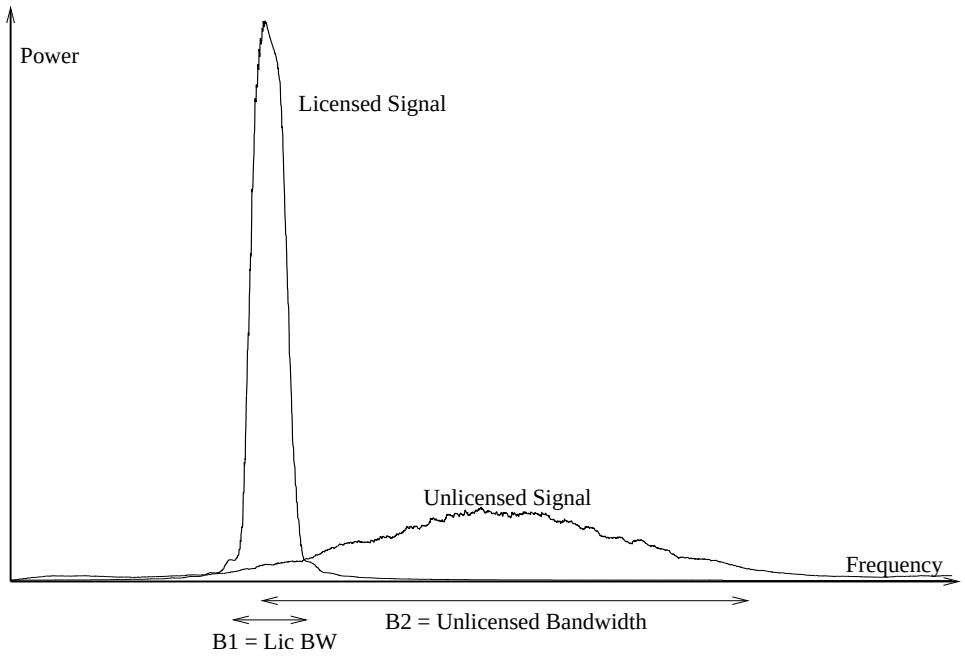

Figure 3.1: Example PSD for an unlicensed signal partially overlapping a licensed signal

pletely characterize both interference and noise with a single number. Of course, it has been argued that interference and noise behave differently. Interference is typically more deterministic and uncorrelated to bandwidth, whereas noise is not.

For a given geographic area, the FCC would establish an *interference temperature limit*, T_L. This value would be a maximum amount of tolerable interference for a given frequency band in a particular location. Any unlicensed transmitter utilizing this band must guarantee that their transmissions added to the existing interference must not exceed the interference temperature limit at a licensed receiver.

While this may seem clear cut, there is ambiguity over which signals are considered interference, and which f_c and B to use. Should they reflect the unlicensed transceiver or the licensed receiver? For example, consider Figure 3.1. Should we use B_1 or B_2 as our bandwidth for our computations? These ambiguities precipitate

Table 3.1: Comparison of the ideal and generalized interference temperature models.

Ideal Model
• Interference to unlicensed transmitters included in the interference temperature measurement is defined as background interference and tranmissions from secondary users
• Interference temperature limit is defined over the bandwidth of the licensed signal
Generalized Model
• Interference to unlicensed transmitters included in the interference temperature measurement is defined as background interference, transmissions from primary users, and tranmissions from secondary users
• Interference temperature limit is defined using the bandwidth of the unlicensed transmission

the need for the two interpretations described in Table 3.1.

3.1.1 Ideal Model

In the *ideal interference temperature model* we attempt to limit interference specifically to licensed signals. Assume our unlicensed transmitter is operating with average power P, and frequency f_c, with bandwidth B. Assume also that this band $[f_c - B/2, f_c + B/2]$ overlaps n licensed signals, with respective frequencies and bandwidths of f_i and B_i. Our goal is to then guarantee that

$$T_I(f_i, B_i) + \frac{M_i P}{k B_i} \leq T_L(f_i) \quad \forall\ 1 \leq i \leq n \tag{3.2}$$

In other words, we guarantee that our transmission does not violate the interference temperature limit at licensed receivers.

Note the introduction of constants M_i. This is a fractional value between 0 and 1, representing a multiplicative attenuation due to fading and path loss between the unlicensed transmitter and the licensed receiver. The idea is that the interference temperature model restricts interference at the licensed receiver, not the unlicensed transmitter, and therefore we must account for attenuation between these two devices. Since we cannot know our distance to all licensed receivers, let us assume that this value is fixed by a regulatory body to a single constant M.

There are two main challenges in implementing the ideal model. The first involves identifying licensed signals. One key question arises: how do you distinguish licensed signals from unlicensed ones? For specific cases, this can be relatively easy. In particular, consider the problems faced by IEEE 802.22 [13], currently under investigation by their spectrum sensing task group. They wish to coexist with digital television (DTV) signals, and can implement very specialized, matched filter sensors to look for DTV transmitters. If you know exactly with whom you are coexisting, then this problem becomes simpler.

The second problem involves measuring T_I in the presence of a licensed signal. We wish to measure the interference floor *underneath* the licensed signal. Again, this can be relatively easy if we have knowledge of the licensed waveform's structure. For example, with DTV, we can measure during the blanking interval when the signal is not present. Also, if we have precise knowledge of the signal's bandwidth B and center frequency f_c, we can approximate the interference temperature as

$$T_I(f_c, B) \approx \frac{P(f_c - B/2 - \tau) + P(f_c + B/2 + \tau)}{2kB} \tag{3.3}$$

where $P(f)$ is the sensed signal power at frequency f and τ is a safety margin of a few kHz.

Assuming a specialized environment where we can locate licensed signals and measure interference temperature, our next goal is to determine radio parameters f_c, B, and P that achieve a desired capacity C. This will be a piecewise-continuous optimization problem with constraints defined in (3.2). If we use a sculptable waveform like OFDM we may be able to more easily meet the various constraints.

On interesting problem related to this model is what happens if you don't overlap any licensed signals, or the signals are so low power that we cannot detect them? There would be no maximum power constraint, and if there were undetected signals, we could cause harmful interference.

We could impose the interference temperature limit across the entire frequency band, rather than just where licensed signals are detected. The problem with this is that no bandwidth is defined over which to apply the interference temperature limit. A regulatory body could fix a bandwidth, but then the interference temperature limit becomes a simple power constraint.

3.1.2 Generalized Model

The generalized interference temperature model, on the other hand, has a different interpretation to signals and bandwidths. The fundamental premise of the generalized model is that we have no a priori knowledge of our signal environment, and consequently have no way of distinguishing licensed signals from interference

and noise.

Under these assumptions, we must apply the interference temperature model to the entire frequency range, and not just where licensed signals are detected. This translates into the following constraint.

$$T_I(f_c, B) + \frac{MP}{kB} \leq T_L(f_c) \tag{3.4}$$

Notice that the constraint is in terms of the *unlicensed transmitter's* parameters, since the parameters of the licensed receivers are unknown. One question that immediately comes to mind: under what conditions does the generalized model limit interference as well as the ideal model?

If we rewrite our constraints in terms of P and combine them, we obtain the following requirement:

If we solve both constraint equations (3.2) and (3.4) for P, we obtain the following equations

$$\begin{aligned} P^{id} &= B_i(T_L(f_c) - T_I^{id}(f_i, B_i)) \\ P^{gen} &= B(T_L(f_c) - T_I^{gen}(f_c, B)) \end{aligned} \tag{3.5}$$

To cause *less* interference in the generalized, we are interested in the case where $P^{id} \geq P^{gen}$. Combining, we obtain the following relation

$$\begin{aligned} B(T_L(f_c) - T_I^{gen}(f_c, B)) \leq B_i(T_L(f_c) - T_I^{id}(f_i, B_i)) \\ \forall\ 1 \leq i \leq n \end{aligned} \tag{3.6}$$

Assuming each licensed signal has power P_i and otherwise the interference floor is defined by the thermal noise temperature T_N, we can transform (3.6) into the following:

$$kBT_L(f_c)(B - B_i) + kBT_N \sum_{j=1}^{n} B_j \leq \sum_{j=1}^{n} B_j P_j \qquad \forall\ 1 \leq i \leq n \tag{3.7}$$

In general, provided B_i and P_i are sufficiently large, this condition can be easily met.

If we consider only one licensed receiver, the inequality simplifies to

$$\frac{kBT_L}{P_1 - kBT_N} \leq \frac{B_1}{B - B_1} \tag{3.8}$$

Thus a small T_L, large B_1, or large P_1 will generally satisfy the constraint.

The next section describes challenges inherent in selecting transmission bandwidths necessary to meet a particular target capacity in each interference temperature model.

3.2 Measuring Interference Temperature

In this section we describe techniques for selecting a power and bandwidth to meet a particular capacity requirement. Each model imposes different constraints on how interference temperature (IT) should be measured, and what those measurements signify.

3.2.1 Properties of Interference Temperature

One shortcoming in the design of the interference temperature model is its simplicity. The goal was to define a single metric that fully captures both the properties of interference and noise. In the end, a *temperature* approach was used rather than a *power* approach. This accurately models the noise portion of the metric, but not the interference portion.

Our eventual goal is to determine the difference between the regulatory interference temperature limit and the measured interference temperature. This then defines the *transmission temperature* our cognitive radio can use, where for a given bandwidth we can compute the maximum allowed power.

Let's define things a little more concretely. Thus, the interference temperature T_I can be specified as a function of bandwidth B as

$$\begin{aligned} T_I(f_c, B) &= \frac{1}{Bk} P_I(f_c, B) \\ &= \frac{1}{Bk} \left(\frac{1}{B} \int_{f_c - B/2}^{f_c + B/2} S(f)\, df \right) \\ &= \frac{1}{B^2 k} \int_{f_c - B/2}^{f_c + B/2} S(f)\, df \end{aligned} \tag{3.9}$$

where $S(f)$ represents power spectral density of our current RF environment.

Recall that in the ideal model, P_I reflects only the interference and noise, where in the generalized model, P_I reflects both the interference, noise, and any licensed signals. This same characterization extends to $S(f)$.

Next, we must consider how our transmission will affect the received inter-

ference temperature $\hat{T}_I(f_c, B)$. As described before, the end goal is to compute a transmit power P and bandwidth B that satisfy our constraints (3.2) and (3.4), depending on our model.

There are two basic cases to consider. First, B is known, and we wish to compute a valid P. In the ideal model, this is fairly straightforward. Rewriting (3.2) we have

$$P \leq \frac{B_i k}{M}(T_L(f_i) - T_I(f_i, B_i)) \quad \forall\ 1 \leq i \leq n \tag{3.10}$$

where the assumption is that for a selected B we overlap n licensed signals with parameters f_i and B_i respectively. If $n = 0$ then we must have $P \leq P_{\max}$, the radio's maximum transmit power. For $n > 0$, to meet this constraint, we minimize of i:

$$P \leq \min_{i \in [1..n]} \left(\frac{B_i k}{M}(T_L(f_i) - T_I(f_i, B_i)) \right) \tag{3.11}$$

This gives us a way to compute P as a function of B.

In the generalized model, we can solve for P and get

$$P \leq \frac{Bk}{M} T_L(f_c) - \frac{1}{BM} \int_{f_c - B/2}^{f_c + B/2} S(f)\, df \tag{3.12}$$

If we are trying to compute a valid B in terms of P, things get a little more complicated for both models. In the ideal model, our maximum bandwidth is going to depend on the interference at various licensed signals. Let there be n^* signals possibly overlappable by our radio for a given f_c and maximum transmit bandwidth $B_{\max}$. For each, licensed signal we can compute our maximum transmit power in

that band P_i as

$$P_i = \frac{B_i k}{M}(T_L(f_i) - T_I(f_i, B_i)) \tag{3.13}$$

If $\forall\ 1 \leq i \leq n^*$, $P_i > P$, then we will not cause any harmful interference regardless of our bandwidth, and therefore our constraint is

$$B \leq B_{\max} \tag{3.14}$$

Otherwise, we can find the index i^* of the signal closest to f_c to which we will cause harmful interference as

$$i^* = \underset{i \in \{i: P_i < P\}}{\arg\max} |f_c - f_i| \tag{3.15}$$

From this we can compute the maximum bandwidth as

$$B \leq 2\left(|f_c - f_{i^*}| - B_{i^*}/2\right) \tag{3.16}$$

For the generalized model, there is no closed-form solution for a general $S(f)$. However, since $S(f)$ is a real, nondecreasing, continuous function of B, there is a solution $\forall S(f)$, even though it may be outside our radio's dynamic range of $(0, B_{\max}]$.

Consider the degenerate case where $S(f) = c$, where c is much larger than the noise floor. This implies some constant level of interference throughout our frequency band of interest. The following solution then arises

$$B \geq \frac{MP + c}{T_L(f_c)k} \tag{3.17}$$

Interestingly, this now indicates a *minimum* bandwidth required for our transmission, not a maximum. This is because of how the interference temperature limit works in the generalized model. Our maximum transmit power is BT_Lk, which increases as a function of bandwidth. This implies that as we use more spectral resources in the frequency domain, we can actually cause *more* interference. This subtlety is counterintuitive, but it is beneficial for our radios, since an increasing noise floor can be offset by an increasing power constraint.

Since bandwidth and power are so interrelated, in the next sections we consider them jointly in terms of capacity.

3.2.2 Capacity in the Ideal Model

So far we've bounded bandwidth in terms of power, and vice versa. Let's change the formulation somewhat, and consider them jointly in terms of capacity. The Shannon-Hartley Theorem states

$$C = B \log_2 \left(1 + \frac{LP}{P_I + P_L} \right) \tag{3.18}$$

where B and P are as before, P_I represents interference power, and P_L represents the average power contributed by licensed signals.

Notice the addition of another constant, L. This value is similar to M, except it represents multiplicative path loss between the unlicensed transmitter and unlicensed receiver. We are measuring capacity at the *receiver*, and therefore need knowledge of the bandwidth and power at the receiver.

As before, let n be the number of licensed signals we overlap[1]. However, let it be a function of f_c and B, such that $n(f_c, B)$ is the number of signals we overlap the frequency range $[f_c - B/2, f_c + B/2]$.

Let's assume we have a bandwidth B. The maximum transmit power is

$$P^*(f_c, B) = \begin{cases} P_{\max} & n(f_c, B) = 0 \\ \min\left(P_{\max}, \min_{i \in [1..n(f_c,B)]} \left(\frac{B_i k}{M}(T_L(f_i) - T_I(f_i, B_i))\right)\right) & n(f_c, B) > 0 \end{cases} \tag{3.19}$$

Note that $P^*(f_c, B)$ is a non-increasing function of B. As we increase our bandwidth, we overlap more signals that could lower our transmission power.

Looking at P_I and P_L, we can compute the interference to our transmission as

$$\begin{aligned} P_I(f_c, B) &= kBT_I(f_c, B) \\ P_L(f_c, B) &= \frac{1}{B} \sum_{i=1}^{n(f_c,B)} P_i B_i \end{aligned} \tag{3.20}$$

Interference $P_I(f_c, B)$ is increasing with B, as the noise floor increases due to thermal noise. We cannot say anything about $P_L(f_c, B)$: it could either be increasing, decreasing, or both.

Thus our achievable capacity is

$$C^*_{id}(f_c, B) = B \log_2 \left(1 + \frac{LP^*(f_c, B)}{P_I(f_c, B) + P_L(f_c, B)}\right) \tag{3.21}$$

[1]For simplicity, we do not examine partially overlapping signals. The analysis could be extended to account for this, but the notation becomes particularly awkward. Capacity would then become a continuous function of B.

As long as $P^*/(P_I + P_L)$ is decreasing at sub-exponential rate, increasing B will generally increase C. However, it will be highly dependent on the RF environment.

In a real radio, assuming the signal processing issues associated with the ideal model are solved, computing P and B subject to some C should be relatively simple. For a given f_c, simply characterize all n^* licensed signals and measure the interference temperature at each. From that data, a numeric version of $C^*_{id}(f_c, B)$ can be calculated, and solved for C.

3.2.3 Capacity in the Generalized Model

Interference temperature must always be measured at some bandwidth B, due to deterministic interference sources that are accentuated in the generalized model. To measure $T_I(f_c, B)$, down-sample the passband signal such that f_c is at $B/2$. Then quantize the spectrum at rate $2B$, and compute its power spectral density (PSD). This will yield a power spectrum for the frequency range $f_c - B/2$ to $f_c + B/2$, which is $\hat{S}_B(f)$. To compute the interference temperature, integrate as follows.

$$T_I(f_c, B) = \frac{1}{B^2 k} \int_0^B \hat{S}_B(f)\, df \tag{3.22}$$

Thus, we can now compute our interference temperature as a function of B.

Let's say a minimum capacity of C is necessary for our communication. The next goal is to find a P and B that both meet regulatory requirements and achieve our capacity constraints. In the last section, we showed that choosing a B and solving for a maximum P was a simple approach. However, considering $S(f)$ may

have very steep slopes when f_c is close to a powerful licensed signal, this may be problematic.

Thus, we must combine some of our concepts. We must compute a capacity function $C^*_{gen}(f_c, B)$ in terms of B, and then solve $C^*_{gen}(f_c, B) = C$ for B. Let's assume a maximum transmit power is used for our bandwidth selection, or

$$P^*(f_c, B) = \frac{Bk}{M}(T_L(f_c) - T_I(f_c, B)) \tag{3.23}$$

This equation uses $T_I(f_c, B)$, measured using the technique described above.

Then we can define our capacity as

$$\begin{aligned} C^*_{gen}(f_c, B) &= B \log_2 \left(1 + \frac{LBk(T_L(f_c) - T_I(f_c, B))}{MBkT_I(f_c, B)}\right) \\ &= B \log_2 \left(1 + \frac{L(T_L(f_c) - T_I(f_c, B))}{MT_I(f_c, B)}\right) \end{aligned} \tag{3.24}$$

This result has similar properties as capacity in the ideal model. Generally, it is increasing with B, but could vary greatly from RF environment to RF environment.

3.2.4 Solving for Capacity

Solving $C^*(f_c, B) = C$ can be difficult for both models[2]. For a general interference environment, this must be done numerically. Figure 3.2 shows a simple example of a 10 MHz licensed signal with square power spectral density located 10 MHz from our carrier frequency. We can see that as long as the signal's power is relatively low, e.g. -90 dBm, the capacity function for the generalized model remains

[2]We use C^* to represent both C^*_{id} and C^*_{gen}, when distinguishing is unimportant.

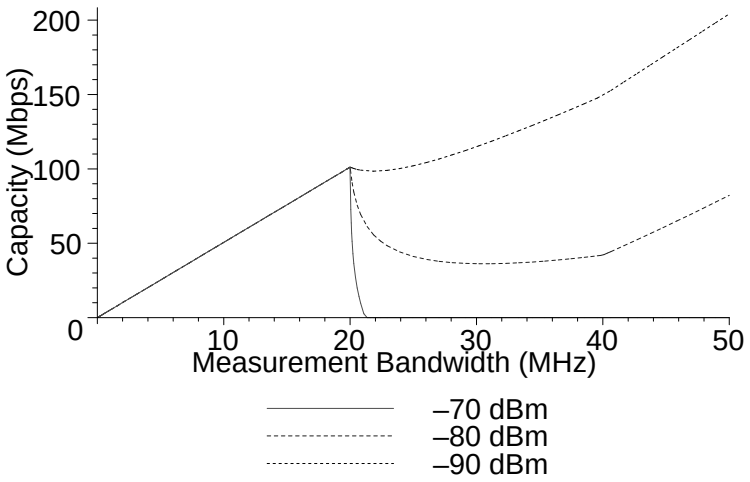

Figure 3.2: Example capacities as a function of B for the generalized model, assuming a licensed signal of varying strengths located at $[f_c + 10 \text{ MHz}, f_c + 20 \text{ MHz}]$, with $T_L = 10000$ Kelvin and a noise temperature of 300 Kelvin. As the interference power increases, the capacity past 20 MHz falls off more.

relatively linear. However, for -70 dBm, we can see that it will significantly hamper our capacity. We can achieve maximum capacity if we avoid the signal all together.

As the previous example illustrates, the capacity function is not strictly increasing, and therefore there may be multiple bandwidths that give the same capacity. Certainly the best choice is to select the smallest bandwidth possible that will achieve your desired capacity.

One could even take that a step further and add a pricing function. In the previous example, if the interference power is -80 dBm, to go from a capacity of 100 Mbps and a capacity of 105 Mbps requires a tripling in the bandwidth. A pricing function would penalize nodes who use extremely large bandwidths, and therefore select 20 MHz, even if it didn't completely satisfy its capacity constraints, as the payoff for tripling the bandwidth would not be worth the added cost.

Putting the pricing function aside, and assuming we have a hard capacity

constraint, and we wish to solve $C^*(f_c, B) = C$ for B, then we must employ numeric techniques. Using the above equations, for a particular B we can compute $T_I(f_c, B)$ and consequently $C^*(f_c, B)$.

Hill Climbing Approach

We can frame the problem as a constrained optimization problem with objective function

$$|C^*(f_c, B) - C| \tag{3.25}$$

One approach is to hill climb, trying to minimize our objective function with respect to B [26]. This function may have several global minimizers over the bandwidth range of our radio. Our goal is to locate the one corresponding to the smallest bandwidth.

A good approach is to run our hill climbing algorithm several times with

$$B_0 = \left\{ \frac{iB_{\max}}{N} \right\}_{i=1..N} \tag{3.26}$$

This will yield N, likely non-unique, solutions. Simply select the one with the smallest bandwidth.

A simple pricing scheme can also be used. To find global capacity maximizers, set $C = \infty$ and run the same algorithm. This will yield a set of points $(B_i, C_i)_{i=1..N}$. Let our capacity utility function be $U(C)$ and our pricing function be $P(B)$. We

can then select the best point i^* as

$$i^* = \arg\max_{i=1..N} \left(U(C_i) - P(B_i)\right) \tag{3.27}$$

Lastly, we should address selection of N. The number of local minima will be proportional to the number of interfering signals. This could be computed by the radio by determining the number of local maxima n in $S(f)$ for $f_c - B/2 \leq f \leq f_c + B/2$. If solving for a specific C, let $N > 2n$, since there would likely be a solution on either side of the signal. If searching for global capacity maximizers, then $N > n$ should be sufficient. This operation could be done infrequently, and would provide a good estimate for N, assuming interfering signals are relatively uniformly spaced over the target spectrum band.

While we can use the hill climbing approach to both optimize $C^*(f_c, B)$ and solve it for a target capacity, we will see in the next section that fixed-point iteration is a more elegant way to solve $C^*(f_c, B)$ for a target capacity. Therefore, hill climbing is most appropriate when trying to maximize capacity.

Fixed-Point Iteration Approach

Many problems that you can solve with hill climbing can also be solved using fixed-point iteration. The basic idea is to take the original problem, $C^*(f_c, B) = C$ and rewrite it as $B_i = f(C^*(f_c, B_{i-1}), C)$, and hope that $\{B_i\}_{i=0...n}$ converges.

Due to the complexity of $C^*_{id}(f_c, B)$, in this section we will only consider application of fixed-point iteration to the generalized model. Next, consider a refor-

mulation of the original problem.

Theorem 1 *The sequence* $\{B_i\}_{i=1..n}$ *where*

$$B_{i+1} = \frac{C}{\log_2\left(1 + \frac{L(T_L(f_c) - T_I(f_c, B_i))}{MT_I(f_c, B_i)}\right)} \tag{3.28}$$

converges linearly to a solution to

$$C = B\log_2\left(1 + \frac{L(T_L(f_c) - T_I(f_c, B_i))}{MT_I(f_c, B_i)}\right) \tag{3.29}$$

as long as

$$B_0 > \frac{2CT_I^*}{T_N}\log_2\left(1 + \frac{L(T_L(f_c) - T_I^*)}{MT_I^*}\right)^{-2} \tag{3.30}$$

where

$$T_I^* = \max_{B\in(0,B_{\max}]} T_I(f_c, B) \tag{3.31}$$

and a solution exists in $B \in (0, B_{\max}]$.

Proof: We're examining our problem in terms of fixed-point approximation. Let

$$g(B) = \frac{C}{\log_2\left(1 + \frac{L(T_L(f_c) - T_I(f_c, B))}{MT_I(f_c, B)}\right)} \tag{3.32}$$

The theory of fixed-point iteration methods dictates that if $B = g(B)$ has at least one solution in some interval $[a, b]$, $g(B)$ is continuous, and $|g'(B)| < 1$ then any starting point in that interval will converge to a solution [2]. Intersect the interval

$[a, b]$ with our feasible interval, $(0, B_{\max}]$. The result is a range for B_0:

$$B_0 \in [a, \min\{b, B_{\max}\}] \tag{3.33}$$

$T_I(f_c, B)$ is continuous, so consequently $g(B)$ is continuous. The derivative constraint can be expressed as follows:

$$\frac{CLT_L(f_c)|T_I'(f_c, B)|}{T_I(f_c, B)(LT_L(f_c) + (M - L)T_I(f_c, B))} < \log_2\left(1 + \frac{L(T_L(f_c) - T_I(f_c, B))}{MT_I(f_c, B)}\right)^2 \tag{3.34}$$

Obviously this constraint is not entirely useful, as it is in terms of B, which we do not yet know. In order to simply this further, we need to remove our dependence on B. First, we use the definition of T_I^* provided in the theorem statement, and notice that

$$T_N \leq T_I(f_c, B) \leq T_I^* \tag{3.35}$$

Next we need to examine the derivative of our interference temperature.

$$T_I'(f_c, B) = \frac{\hat{S}(B)}{B^2 k} - \frac{2}{B}T_I(f_c, B) \tag{3.36}$$

Thus to maximize $|T_I'(f_c, B)|$, let $\hat{S}(B) = 0$ and $T_I(f_c, B) = T_I^*$. The result is

$$|T_I'(f_c, B)| \leq 2T_I^*/B \tag{3.37}$$

Substituting, we have

$$\begin{aligned} B_0 &> \frac{CLT_L(f_c)2T_I^*}{T_N(LT_L(f_c)+(M-L)T_N)} \log_2\left(1+\frac{L(T_L(f_c)-T_I^*)}{MT_I^*}\right)^{-2} \\ &> \frac{CLT_L(f_c)2T_I^*}{T_N(LT_L(f_c))} \log_2\left(1+\frac{L(T_L(f_c)-T_I^*)}{MT_I^*}\right)^{-2} \\ &> \frac{2CT_I^*}{T_N} \log_2\left(1+\frac{L(T_L(f_c)-T_I^*)}{MT_I^*}\right)^{-2} \end{aligned} \tag{3.38}$$

Thus we have proved our theorem. ∎

We now have a viable algorithm for computing the required bandwidth B in terms of desired capacity C. If $B_0 > B_{\max}$, this does not necessarily mean a solution does not exist, since we derived a *sufficient* condition, and not a *necessary* one. If divergence is detected, then the capacity C must be decreased in order to find a solution.

The key point is that fixed-point iteration can find a solution if one exists, but may not always succeed. As a result, it may be useful to implement a hybrid algorithm that first tries fixed-point iteration, and if divergence is detected, switch over to a hill climbing approach. Note that the algorithms can be executed on a PSD snapshot taken with bandwidth $B_{\max}$, and consequently radio sensing resources need not be tied up during algorithm execution. In Chapter 4, ITMA will be discussed, whose data-link layer does bandwidth selection based on a target capacity. It uses fixed-point iterations to find the required bandwidth.

3.3 Frequency Selection

In the previous sections we describe how to select a bandwidth given a center frequency f_c. However, one of the major uses for cognitive radio is to dynamically select your center frequency to exploit spectrum access opportunities.

There are two main schools of thought on dynamic center frequencies. In particular, the ability to change f_c in real time increases higher-layer protocol complexity, since the receiver must know that the transmitter has changed frequency. These competing ideas are related to how radios exchange radio parameters.

The first assumes there is a management or control channel through which radios can coordinate. Devices could indicate the center frequency, waveform, destination, and time of their next transmission. Thus, f_c is something to be optimized and changed in real time.

However, others consider the management channel an unrealistic assumption. In a dense, busy packet network environment, management of the management channel becomes a problem. Also, how can we guarantee the management channel is not causing harmful interference?

In Chapter 4, we propose a logical management channel embedded within the main channel. This, however, assumes a fairly static center frequency.

Here, we look at how to select f_c for optimal performance, and ignore protocol issues for coordination. We simply address how you can select the best f_c at a particular time. The approach is a simple extension of the ideas in the last section.

We defined our capacity functions for each model in the previous sections, and

described techniques to solving

$$C^*(f_c, B) = C \tag{3.39}$$

for B. However, if we assume f_c is no longer fixed, how does that change things?

We advocate selecting an f_c at the beginning to maximize your eventual per-packet capacity, and leaving it fixed unless communication at that frequency becomes impossible. Thus, the optimal center frequency is

$$f_c^* = \max_{f \in [f_{\min}, f_{\max}]} \left(\max_{B \in (0, B_{\max}]} C^*(f, B) \right) \tag{3.40}$$

Maximizing over B ca be done using the hill climbing approach. Assuming the space of frequencies is channelized, then $[f_{\min}..f_{\max}]$ is a discrete set, and the hill climbing can be executed for each f.

Alternatively, we can look at the structure of $C^*(f_c, B)$ in more detail. In particular, in the presence of uniform interference, both capacity functions are maximized when licensed signals are completely avoided. Assume n licensed signals are detected within our radio's overall candidate frequency band. Let each be located at center frequency f_i and have bandwidth B_i. Assume $\{f_i\}_{i=1}^n$ is an ordered set, where

$$f_1 \leq f_2 \leq \cdots \leq f_n \tag{3.41}$$

Our best frequency is going to be half way between the two signals with fur-

thest distance between them. In particular, if

$$i^* = \arg\max_{i=1..n-1} \left(f_{i+1} - \frac{B_{i+1}}{2} \right) - \left(f_i + \frac{B_i}{2} \right) \tag{3.42}$$

then

$$f_c^* = \frac{1}{2} \left(\left(f_{i^*+1} - \frac{B_{i^*+1}}{2} \right) + \left(f_{i^*} + \frac{B_{i^*}}{2} \right) \right) \tag{3.43}$$

Recall, however, that this assumes our interference is uniform. If interference varies some, but not a significant amount, we can adapt our previous optimization somewhat. In particular, if

$$\left| \frac{d}{df} T_I^{id}(f, B) \right| < \epsilon \quad \forall \ f \in [f_{\min}, f_{\max}] \tag{3.44}$$

then we can define our channelization $\{c_i\}_{i=1}^{n-1}$ as

$$c_i = \frac{1}{2} \left(\left(f_{i+1} - \frac{B_{i+1}}{2} \right) + \left(f_i + \frac{B_i}{2} \right) \right) \tag{3.45}$$

and then maximize over our channels to compute

$$f_c^* = \max_{f=c_1..c_{n-1}} \left(\max_{B \in (0, B_{\max}]} C^*(f, B) \right) \tag{3.46}$$

As discussed, center frequency should be selected to promote a radio environment that will maximize our potential capacity. Typically, this involves steering clear of licensed signals, so we use this fact to pick a set of candidate center fre-

quencies. By computing our maximum capacity at each, we can decide which is optimal.

3.4 Network Capacity Analysis Model

In this section we assume a fixed transmit bandwidth that overlaps a single licensed signal. Our goal is to quantify the total network capacity achievable by the underlay network. Notationally, bandwidths B_U and B_L respectively represent our unlicensed and licensed bandwidths. We use the notation $\mathcal{N}(\mu, \sigma^2)$ to indicate a Gaussian random variable with mean μ and variance σ^2. Also, $\mathbf{exp}(\mu)$ indicates an exponentially distributed random variable with mean μ.

3.4.1 Model Geometry

Here we describe some of our model fundamentals that will be used in later sections.

Lemma 1 *Consider a disc of radius R. The distance $\mathcal{D}$ between a point selected with uniform distribution over the area of the disc and the center of the disc has c.d.f.:*

$$\mathbf{P}(\mathcal{D} \leq x) = \begin{cases} 0 & x < 0 \\ x^2/R^2 & 0 \leq x \leq R \\ 1 & x > R \end{cases} \tag{3.47}$$

Proof: The probability that a point is less than distance x from the center is the ratio of the area of a disc with radius x, and the total area of the disc. Thus we can

compute

$$\mathbf{P}(\mathcal{D} \leq x) = \frac{\pi x^2}{\pi R^2} = x^2/R^2 \tag{3.48}$$

The remainder of the expression is to handle edge cases. ∎

Lemma 2 *Let $\mathcal{P}$ be the λ-wavelength power experienced by a receiver at the center of a disc with radius R, from a single transmitter with position uniformly distributed over the disc, with a transmit power distributed* $\mathbf{exp}(\mu)$. *The expected value and variance of $\mathcal{P}$ are*

$$\begin{aligned} \mathbf{E}[\mathcal{P}] &= \frac{\mu\lambda^2 \log R}{8\pi^2 R^2} \\ \mathbf{Var}[\mathcal{P}] &= \frac{\lambda^4 \mu^2}{128\pi^4 R^4} \left(R^2 - \log^2 R^2 - 1\right) \end{aligned} \tag{3.49}$$

Proof: Consider a disc with radius R. At the center of the disc is a receiver, and surrounding it are transmitters. If a transmitter's location is uniformly distributed, then its distance to the center $\mathcal{D}$ has distribution computed in Lemma 1.

If the transmitted power $\mathcal{T}$ of a signal with wavelength λ has distribution $\mathcal{T} \sim \mathbf{exp}(\mu)$ and experiences path loss[3] over distance $\mathcal{D}$, the received power is

$$\mathcal{P} = \frac{\lambda^2}{16\pi^2 \mathcal{D}^2} \mathcal{T} \tag{3.50}$$

[3]For the purposes of this chapter, we assume a path loss constant of 2, indicating simple free-space path loss. Typically, this value is larger, between 3 and 4, due to the effects of multipath fading. However, using any value other than 2 makes the integrals symbolically uncomputable. These model assumptions must be taken into account when evaluating the results of the analysis based on these models.

This power $\mathcal{P}$ is a random variable defined in terms of random variables $\mathcal{T}$ and $\mathcal{D}$. We can compute its distribution precisely as

$$\begin{aligned}\mathbf{P}(\mathcal{T} \leq x) &= \int_{r_1}^{r_2} \frac{2d}{R^2} \int_0^{16\pi^2 x d^2/\lambda^2} \frac{1}{\mu} \mathbf{e}^{-p/\mu} \, dp \, dd \\ &= \frac{r_2^2 - r_1^2}{R^2} + \frac{1}{\alpha R^2 x} \left(\mathbf{e}^{-\alpha x r_2^2} - \mathbf{e}^{-\alpha x r_1^2} \right)\end{aligned} \tag{3.51}$$

where

$$\alpha = \frac{16\pi^2}{\lambda^2 \mu} \tag{3.52}$$

Notice that we left distance integration limits as r_1 and r_2. If we want to consider transmitters located across the entire disc, we should use $r_1 = 0$ and $r_2 = R$. The latter is fine, however the former causes problems with the laws of physics. In particular, we are using free-space path loss which decays as a function of distance squared. At zero distance a division by zero results.

To work around this problem, we let $r_1 = 1$. This physically corresponds to a guarantee that no transmitters will be within a meter of the receiver. Using this assumption, we have

$$\mathbf{P}(\mathcal{P} \leq x) = \frac{R^2 - 1}{R^2} + \frac{1}{\alpha R^2 x} \left(\mathbf{e}^{-\alpha R^2 x} - \mathbf{e}^{-\alpha x} \right) \tag{3.53}$$

using the same value for α.

Let $P_{\mathcal{P}}(x)$ be the p.d.f. for $\mathcal{P}$, and is computed as

$$\begin{aligned} P_{\mathcal{P}}(x) &= \frac{d}{dx}\mathbf{P}(\mathcal{P} \leq x) \\ &= \frac{1}{R^2 x}\left(1 + \frac{1}{x\alpha}\right)\mathbf{e}^{-x\alpha} - \frac{1}{x}\left(1 + \frac{1}{R^2 x\alpha}\right)\mathbf{e}^{-R^2 x\alpha} \end{aligned} \tag{3.54}$$

If we compute the expected value through integration we get

$$\begin{aligned} \mathbf{E}[\mathcal{P}] &= \int_0^\infty x P_{\mathcal{P}}(x)\, dx \\ &= \frac{\mu\lambda^2 \log R}{8\pi^2 R^2} \end{aligned} \tag{3.55}$$

For the variance can compute it as

$$\begin{aligned} \mathbf{E}[\mathcal{P}^2] &= \int_0^\infty x^2 P_{\mathcal{P}}(x)\, dx \\ &= \frac{\mu^2\lambda^4}{128\pi^4 R^4}(R^2 - 1) \end{aligned} \tag{3.56}$$

and then

$$\begin{aligned} \mathbf{Var}[\mathcal{P}] &= \mathbf{E}[\mathcal{P}^2] - \mathbf{E}[\mathcal{P}]^2 \\ &= \frac{\lambda^4\mu^2}{128\pi^4 R^4}\left(R^2 - \log^2 R^2 - 1\right) \end{aligned} \tag{3.57}$$

Thus proving our lemma. ∎

Now, we're going to change the geometry somewhat, and introduce another disc. Consider two concentric discs C_1 and C_2, with radii R_1 and R_2, respectively, with $R_1 \ll R_2$. Assume that C_2 contains RF transmitters uniformly distributed over the area with density δ_2. Assume their transmit power is exponentially distributed

with mean μ_2, and their transmission wavelength is λ.

Theorem 2 *The signal power $\mathcal{P}_2$ from radios in C_2 as seen in C_1 is normally distributed as follows:*

$$\mathcal{P}_2 \sim \mathcal{N}\left(\frac{\lambda^2 \mu_2 \delta_2 \log R_2}{8\pi}, \frac{\lambda^4 \mu_2^2 \delta_2}{128\pi^3 R_2^2}(R_2^2 - \log^2 R^2 - 1)\right) \tag{3.58}$$

Proof: This is simply an application of our previous lemma. There are $\delta_2 \pi R_2^2$ i.i.d. transmitters, so their total power is normally distributed and can be computed using the Central Limit Theorem. The above values result. ∎

This result is particularly interesting. First, notice that as R_2 increases, our mean increases logarithmically. This is an intuitive result, since nodes further away will contribute a diminishing amount to the interference environment. Also intriguing is that the variance is constant with respect to R_2.

The mean being logarithmic allows the large-scale estimation done in Sections 3.4.2 and 3.4.3. As long as $R_1 \ll R_2$, interference effects are roughly constant throughout C_1, since $\log(R_2) \approx \log(R_2 - R_1)$.

Corollary 1 *A reasonable upper bound for $\mathcal{P}_2$ is:*

$$\mathbf{P}\left[\mathcal{P}_2 < \frac{\lambda^2 \mu_2}{8\pi^2 R_2}\left(\delta_2 \pi R_2 \log R_2 + \sqrt{2\delta_2 \pi (R^2 - \log^2 R^2 - 1)}\right)\right] > 0.98 \tag{3.59}$$

Proof: A good confidence interval is $\mu_{\mathcal{P}_2} + 2\sigma_{\mathcal{P}_2}$, which is the value used above. ∎

Next, let's move the transmitters to C_1. Assume C_1 contains RF transmitters, uniformly distributed over the area with density δ_1. Assume their transmit power is exponentially distributed with mean μ_1, and their transmission wavelength is λ.

Theorem 3 *The signal power $\mathcal{P}_1$ from radios in C_1 as seen in C_2 at a distance r from the center with $R_1 \ll r < R_2$ is normally distributed as follows:*

$$\mathcal{P}_1 \sim \mathcal{N}\left(\frac{R_1^2\delta\mu_1\lambda^2}{16\pi r^2}, \frac{R_1^2\delta\mu_1^2\lambda^4}{128\pi^4 r^2}\right) \tag{3.60}$$

Proof: Here we apply the Central Limit Theorem to $\pi\delta R_1^2$ nodes, each with exponentially distributed power at roughly distance r from the receiver. This total power then undergoes free-space path loss, and the above distribution results.

In particular, for a single node transmitting with power $\mathcal{T}$, we have receive power $\mathcal{P}$ where

$$\begin{aligned} \mathcal{P} &= \frac{\lambda^2}{16\pi^2 r^2}\mathcal{T} \\ &\sim \mathbf{exp}\left(\frac{\lambda^2\mu_1}{16\pi^2 r^2}\right) \end{aligned} \tag{3.61}$$

This has moments

$$\begin{aligned} \mathbf{E}[\mathcal{P}] &= \frac{\lambda^2\mu_1}{16\pi^2 r^2} \\ \mathbf{Var}[\mathcal{P}] &= \frac{\lambda^4\mu_1^2}{128\pi^4 r^4} \end{aligned} \tag{3.62}$$

We use the Central Limit Theorem to sum all transmitters. The result is as specified above. ∎

Next we're going to use Lemma 1 to prove some minimum distance bounds that will be used later.

Lemma 3 *Let $\delta\pi R^2$ points be randomly placed over a disc of radius R, with density δ. Let $\mathcal{D}_{\min}$ be a random variable representing the distance between the center of the disc, and the point closest to the center of the disc. The c.d.f. of $\mathcal{D}_{\min}$ is*

$$\mathbf{P}(\mathcal{D}_{\min} < d) = 1 - \mathbf{e}^{-d^2\pi R\delta/2} \tag{3.63}$$

Proof: The distance of each of point and the center of the disc a random variable $\mathcal{D}_i$, as defined by Lemma 1. Our goal is to determine the distance distribution for the closest one.

$$\mathcal{D}_{\min} = \min_{i=0..N_L} \mathcal{D}_i \tag{3.64}$$

The resulting distribution for $\mathcal{D}_{\min}$ is the Rayleigh distribution [9].

$$\begin{aligned} P_{\mathcal{D}_{\min}}(x) &= \mathbf{Rayleigh}(R/N_W, x) \\ &= \mathbf{Rayleigh}(1/\delta_W\pi R, x) \\ &= x\delta_W\pi R\mathbf{e}^{-x^2\delta_W\pi R/2} \end{aligned} \tag{3.65}$$

We can compute the c.d.f. by integrating, and obtain

$$\mathbf{P}(\mathcal{D}_{\min} < d) = 1 - \mathbf{e}^{-d^2\pi R\delta/2} \tag{3.66}$$

Thus, we have proved the lemma. ∎

Next, let's define the idea of density uniformity. In particular, if we say area A has node density δ, then that means we have a total of δA nodes in area A. However, this could imply that all nodes are located in a single corner of A, and when looking at some area $A' < A$, we could discover a different node density. While our original density was correct for A, it is no longer correct for A'.

Let's define our density in terms of the area, $\delta(A)$. Density uniformity defines a minimum area $A_{\min}$ for which

$$\mathbf{P}(|\delta(A_{\min}) - \delta(A)| > \epsilon_1) < \epsilon_2 \tag{3.67}$$

for some tolerances ϵ_1 and ϵ_2.

Corollary 2 *Given density δ and minimum area $\pi R_{\min}^2$, with probability p we can be sure the distance between a point and its closest neighbor is at least d is*

$$d = \frac{\sqrt{-2\log(p)}}{\pi R_{\min}\delta} \tag{3.68}$$

Proof: Application of the previous lemma to the described scenario. ∎

These theorems will provide the foundation for the computations performed in the upcoming sections.

3.4.2 Wireless WAN

In this section, we define a wireless WAN (WWAN) to be a wireless network utilizing the interference temperature model that has an operational radius signifi-

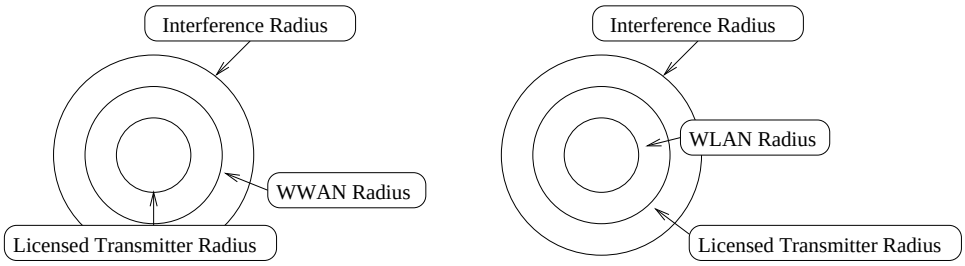

Figure 3.3: Diagram of the two models. In the first, we have a wireless wide-area network (WAN), in which radii are as follows: $R_L \ll R_W \ll R_I$. In the second, the wireless local area network (LAN) has radii $R_W \ll R_L \ll R_I$.

cantly larger than our licensed radio network. Figure 3.3 illustrates this. An example of this could be a mobile broadband radio network covering hundreds of kilometers, coexisting with a traditional UHF TV broadcasting station.

The goal of our analysis is to examine the capacity that can be achieved by an underlay network. Our first step is to determine the base interference temperature seen by our WAN nodes.

Using Theorem 2, we can compute the distribution on $\mathcal{P}_I$ in terms of our average interferer power μ_I, interferer density δ_I, and radius R_I. However, since interference temperature is likely to change frequently, radios should measure it over some fixed time period and use the maximum recorded value. This will correspond to our confidence interval value from Corollary 1. As a result, we estimate

$$P_I = \frac{\lambda^2 \mu_I}{8\pi^2 R_I} \left(\delta_I \pi R_I \log R_I + \sqrt{2\delta_I \pi (R_I^2 - \log^2 R_I^2 - 1)} \right) \tag{3.69}$$

Now assume that our WAN node transmit power is exponentially distributed with mean μ_W. This mean is going to be a function of our interference temperature

and interference temperature limit. The interference temperature is going to include power from interferers P_I, unlicensed transmitters P_W, and in the generalized model also the licensed signal P_L. The power of the interferers was just determined as P_I. The power from the unlicensed devices must be represented in terms of μ_W.

Computing a probability distribution for $\mathcal{P}_W$ yields some rather nasty integrals, and as a result we will use the worst case of the transmitter being located at the center of the unlicensed network. Thus we can use reuse Corollary 1.

$$\begin{aligned} P_W &= \mathbf{E}[\mathcal{P}_W] \\ &\leq \frac{\lambda^2 \mu_W}{8\pi^2 R_W} \left(\delta_W \pi R_W \log R_W + \sqrt{2\delta_W \pi (R_W^2 - \log^2 R_W^2 - 1)} \right) \end{aligned} \tag{3.70}$$

Use the following value for the ideal model

$$\begin{aligned} \mu_W^{id} &= \frac{T_L k B_L}{M} - P_I - P_W \\ &= \frac{T_L k B_L}{M} - \frac{\lambda^2 \mu_I}{8\pi^2 R_I} \left(\delta_I \pi R_I \log R_I + \sqrt{2\delta_I \pi (R_I^2 - \log^2 R_I^2 - 1)} \right) \\ &\quad - \frac{\lambda^2 \mu_W^{id}}{8\pi^2 R_W} \left(\delta_W \pi R_W \log R_W + \sqrt{2\delta_W \pi (R_W^2 - \log^2 R_W^2 - 1)} \right) \end{aligned} \tag{3.71}$$

Solving for μ_W^{id}, we have

$$\mu_W^{id} = \frac{8\pi^2 T_L k B_L / M - \lambda^2 \mu_I \left(\delta_I \pi \log R_I + R_I^{-1} \sqrt{2\delta_I \pi (R_I^2 - \log^2 R_I^2 - 1)} \right)}{8\pi^2 + \lambda^2 \left(\delta_W \pi \log R_W + R_W^{-1} \sqrt{2\delta_W \pi (R_W^2 - \log^2 R_W^2 - 1)} \right)} \tag{3.72}$$

For the generalized model, we must also subtract the average power from the licensed transmitter P_L, and evaluate T_L over B_U rather than B_L. Assuming the

licensed signal's power is $\mathbf{exp}(\mu_L)$, we can compute P_L using Lemma 2.

$$\begin{aligned} P_L &= \mathbf{E}[\mathcal{P}_L] \\ &= \frac{\mu_L \lambda^2 \log R_W}{8\pi^2 R_W^2} \end{aligned} \tag{3.73}$$

Substituting into μ_W^{gen}, we obtain the relation

$$\begin{aligned} \mu_W^{gen} &= \frac{T_L k B_U}{M} - P_I - P_W - \frac{B_L}{B_U} P_L \\ &= \frac{T_L k B_L}{M} - \frac{\lambda^2 \mu_I}{8\pi^2 R_I} \left(\delta_I \pi R_I \log R_I + \sqrt{2\delta_I \pi (R_I^2 - \log^2 R_I^2 - 1)} \right) \\ &\quad - \frac{\lambda^2 \mu_W^{id}}{8\pi^2 R_W} \left(\delta_W \pi R_W \log R_W + \sqrt{2\delta_W \pi (R_W^2 - \log^2 R_W^2 - 1)} \right) \\ &\quad - \frac{B_L \mu_L \lambda^2 \log R_W}{8\pi^2 B_U R_W^2} \end{aligned} \tag{3.74}$$

Again, solving for μ_W^{gen}, we get

$$\mu_W^{gen} = \frac{8\pi^2 T_L k B_L / M - \lambda^2 \left(\mu_I \left(\delta_I \pi \log R_I + R_I^{-1} \sqrt{2\delta_I \pi (R_I^2 - \log^2 R_I^2 - 1)} \right) + \frac{B_L \mu_L \log R_W}{B_U R_W^2} \right)}{8\pi^2 + \lambda^2 \left(\delta_W \pi \log R_W + R_W^{-1} \sqrt{2\delta_W \pi (R_W^2 - \log^2 R_W^2 - 1)} \right)} \tag{3.75}$$

Next, we examine the WAN network capacity, which is defined as the sum of the per-link capacities. Our mean transmit power is μ_W, at bandwidth B_U.

We can then use free-space path loss to compute the received SIR in a CDMA-based access network, averaging the licensed power over the unlicensed bandwidth. This averaging is the appropriate thing to do for CDMA, since our despread operation will effectively spread the narrow-band noise, raising the noise floor for our

despread signal.

$$\mathbf{E}[SIR] = \frac{\mu_W}{\mathbf{E}[\mathcal{P}_I] + \mathbf{E}[\mathcal{P}_W] + \frac{B_L}{B_I}\mathbf{E}[\mathcal{P}_L]} \cdot \frac{\lambda^2}{16\pi^2\mathbf{E}[d^2]} \tag{3.76}$$

Computation of $\mathbf{E}[d^2]$ is a little tricky. The random variable d represents the distance between an unlicensed node and its closest neighbor. We can compute this using the Rayleigh distribution we derived in Corollary 2. In particular, for our density uniformity parameter $R_{\min}$, we know that

$$\mathbf{E}[d^2] = \frac{2}{\pi^2 R_{\min}^2 \delta_W^2} \tag{3.77}$$

Plugging this into our SIR, we obtain,

$$\begin{aligned}\mathbf{E}[SIR] &= \frac{\mu_W \delta_W^2 \lambda^2 R_{\min}^2}{32\pi}\left(\mathbf{E}[\mathcal{P}_I] + \mathbf{E}[\mathcal{P}_W] + \frac{\mathbf{E}[\mathcal{P}_L]B_L}{B_U}\right)^{-1} \\ &= \frac{\pi}{4}\mu_W \delta_W^2 R_{\min}^2 \left(\mu_I \delta_I \log R_I + \mu_W \delta_W \log R_W + \frac{\mu_L B_L \log R_W}{\pi B_U R_W^2}\right)^{-1}\end{aligned} \tag{3.78}$$

Thus our per-link capacity is defined by the Shannon-Hartley Theorem.

$$C = B_U \log\left(1 + \frac{\pi}{4}\mu_W \delta_W^2 R_{\min}^2 \left(\mu_I \delta_I \log R_I + \mu_W \delta_W \log R_W + \frac{\mu_L B_L \log R_W}{\pi B_U R_W^2}\right)^{-1}\right) \tag{3.79}$$

As before, μ_W is defined as μ_W^{id} or μ_W^{gen}, depending on the model. The total network capacity is the per-link capacity multiplied by the number of nodes.

$$C_N = C\pi R_W^2 \delta_W \tag{3.80}$$

For a given set of μ, δ, and R parameters, our per-device capacity is $\mathcal{O}(B_U \log B_U)$.

3.4.3 Wireless LAN

For our Wireless LAN model, we now assume that the radius of our underlay network is significantly smaller than, and wholly contained within our licensed network.

The first steps are completed as before. The computation of P_I and μ_W are the same. The value of μ_W reflects either μ_W^{id} or μ_W^{gen}, as derived in the previous section, with the exception that the wireless network does not have to be at the center of the licensed network, which affects the generalized model. Let r_W, where $0 < r_W < R_L - R_W$ and $R_W \ll R_L$, reflect this distance. Then,

$$\begin{aligned}
\mu_W^{gen} &= \frac{T_L k B_U}{M} - P_I - P_W - \frac{B_L}{B_U} P_L \\
&= \frac{T_L k B_L}{M} - \frac{\lambda^2 \mu_I}{8\pi^2 R_I} \left(\delta_I \pi R_I \log R_I + \sqrt{2\delta_I \pi (R_I^2 - \log^2 R_I^2 - 1)} \right) \\
&\quad - \frac{\lambda^2 \mu_W^{id}}{8\pi^2 R_W} \left(\delta_W \pi R_W \log R_W + \sqrt{2\delta_W \pi (R_W^2 - \log^2 R_W^2 - 1)} \right) \\
&\quad - \frac{B_L \lambda^2 \mu_L}{16\pi^2 B_U r_W^2}
\end{aligned} \tag{3.81}$$

Solving for μ_W^{gen}, we get

$$\mu_W^{gen} = \frac{8\pi^2 T_L k B_L / M - \lambda^2 \left(\mu_I \left(\delta_I \pi \log R_I + R_I^{-1} \sqrt{2\delta_I \pi (R_I^2 - \log^2 R_I^2 - 1)} \right) + \frac{B_L \mu_L}{2 B_U r_W^2} \right)}{8\pi^2 + \lambda^2 \left(\delta_W \pi \log R_W + R_W^{-1} \sqrt{2\delta_W \pi (R_W^2 - \log^2 R_W^2 - 1)} \right)} \tag{3.82}$$

The capacity computations are similar, except the noise from the licensed node

has changed, and is a function of r_W.

$$C = B_U \log \left(1 + \frac{\pi}{4} \mu_W \delta_W^2 R_{\min}^2 \left(\mu_I \delta_I \log R_I + \mu_W \delta_W \log R_W + \frac{B_L \mu_L}{2 B_U r_W^2} \right)^{-1} \right) \tag{3.83}$$

3.5 Impact to Licensed Users

The analysis so far has assumed an interference temperature limit T_L and distance coefficient M have been specified by a regulatory body. However, we have not yet investigated how their values affect licensed transmitters.

For each we derive Δ_{T_L} and Δ_M, which reflect the lower bound for the fractional decrease in coverage area due to the implementation of the interference temperature model. Their cumulative effect, $\Delta_{T_L} \Delta_M$ reflects the total lower bound for decrease in coverage area.

For example, if the original coverage area was 100 square kilometers, and $\Delta_{T_L} = \Delta_M = 0.9$, then the resulting coverage area would be 81 square kilometers *or larger*.

3.5.1 Selection of T_L

Consider a licensed signal at frequency f_c using bandwidth B_L. For most cases, the generalized model will cause less interference than the ideal model, so here we focus on the ideal model.

Provided the interference temperature limit is met at all licensed receivers, in

a worst-case scenario the noise floor will move from $T_I(f_c, B)kB_L$ to $T_L kB_L$ at each receiver. Let $\bar{T}_I$ be the average interference temperature measured over the entire area of our licensed receivers.

Let's assume the same SIR is required to receive signals in both the original environment and the new environment where the interference temperature model is employed. This means the following relationship is true, based on path loss:

$$\frac{\mu_L \lambda^2}{kB_L T_L 16\pi^2 R_L'^{\alpha}} = \frac{\mu_L \lambda^2}{kB_L \bar{T}_I 16\pi^2 R_L^{\alpha}} \tag{3.84}$$

where $\alpha = 2$ represents standard free-space path loss.

Here R_L was the original licensed signal range, and R_L' is the new signal range. We can cancel many of the variables, and the resulting relation results

$$T_L = \frac{\bar{T}_I}{\Delta_{T_L}^{\alpha/2}} \tag{3.85}$$

where Δ_{T_L} represents the fraction of the original coverage area remaining once the interference temperature model has been established. We can easily rewrite in terms of Δ_{T_L} as

$$\Delta_{T_L} = \left(\frac{\bar{T}_I}{T_L}\right)^{2/\alpha} \tag{3.86}$$

Interestingly, or fractional increase in noise floor is directly proportional to the fractional decrease in licensed signal coverage area for $\alpha = 2$.

According to the original FCC specification [10], a likely interference temperature limit would be $(\max T_I)$ over some time period, thus allowing unlicensed

transmission in the existing interference. We can compute this as a confidence interval on $\mathcal{P}_I$. Define the interference temperature limit as

$$T_L = \mathbf{E}[\mathcal{P}_I] + \beta\sqrt{\mathbf{Var}[\mathcal{P}_I]} \tag{3.87}$$

where $\beta \geq 2$.

From this we can compute our fractional decrease per our prior derivation.

$$\begin{aligned} \Delta_{T_L} &= \left(\frac{\mathbf{E}[\mathcal{P}_I]}{\beta\sqrt{\mathbf{Var}[\mathcal{P}_I]} \right)^{2/\alpha} \\ &= \left(\frac{\delta_I \pi \log R_I}{\delta\pi \log R_I + \beta\sqrt{\delta_I \pi/2}} \right)^{2/\alpha} \end{aligned} \tag{3.88}$$

Note that this value depends only on the density of interferers and the interferer radius, and not on their power.

3.5.2 Selection of M

The variable M represents the attenuation due to path loss between an unlicensed transceiver and a licensed receiver. Its use is intrinsic to the interference temperature model which dictates interference *received*, not *transmitted*. In this section we assume free-space path loss, since unlicensed devices typically operate over shorter distances where the free-space model is most accurate.

In this section we propose a mathematical formulation based on Corollary 2 for selecting a value of M to minimize the number of receivers to whom we inadvertently cause harmful interference.

Suppose a single unlicensed transmitter is surrounded by licensed receivers that are randomly placed with a uniform distribution. Let the average density of the receivers be δ_L devices per unit area.

From Corollary 2, we can compute the distance d necessary to guarantee that with probability p we are distance d from the closest licensed receiver.

$$d = \frac{1}{\pi R_{\min} \delta_L} \sqrt{-2 \log(p)} \tag{3.89}$$

We can also define our distance d in terms of the fraction decrease Δ_M of devices with no harmful interference as

$$\begin{aligned} d &= \frac{1}{\pi R_{\min} \delta_L} \sqrt{-2 \log\left(\Delta_M^{1/\pi \delta_W R_W^2}\right)} \\ &= \frac{1}{\pi R_{\min} R_W \delta_L} \sqrt{\frac{-2 \log \Delta_M}{\pi \delta_W}} \end{aligned} \tag{3.90}$$

We can convert this to M using free-space path loss

$$\begin{aligned} M &= \frac{\lambda^2}{16 \pi^2 d^2} \\ &= -\frac{\lambda^2 R_{\min}^2 R_W^2 \delta_L^2 \pi \delta_W}{32 \log \Delta_M} \end{aligned} \tag{3.91}$$

Consider an example where 200 television sets exist in a single square kilometer, yielding $\delta_L = 200$. Let $R_{\min} = 50$ meters, equating with us being sure that there are 8 television sets per 4 houses. Imagine we want to form an underlay network with 10 nodes. Using our relations, we can compute $d = 20$ meters to ensure that with 98 percent probability we will cause no harmful interference, or in other words

interfere with 4 televisions on average.

3.6 Conclusion

In this chapter, we've developed algorithms to measure interference temperature and analyzed the interference temperature model from a purely stochastic perspective. We have proved that realizable network capacity is independent of almost all our model parameters. We have also derived how the underlay network will affect the licensed signal, showing a fractional decrease in coverage area equal to $\Delta_M \Delta_{T_L}$ which can be computed from M and T_L.

One major bullet for future work is the model for analyzing network capacity. In particular, we assume free-space path loss, because no closed-form solutions exist for the capacity in more higher-order RF propigation models. This area deserves further study, though we suspect the only practical means of accurate analysis would be through the simulation conducted in Chapter 4.

Chapter 4

Interference Temperature Multiple Access

This chapter describes a new multiple access technique called *interference temperature multiple access* (ITMA). ITMA relies on the cognitive radio's ability to sense its environment and regulate bandwidth and power usage on a per-packet basis. It uses interference temperature to sense its environment, and transmits using the bandwidth and power derived from the analysis in Chapter 3.

In this chapter, we define the physical (PHY) and medium access control (MAC) layers for ITMA. The main goal of the PHY is to support dynamic bandwidths and powers on a per-packet basis. The lower MAC layer is responsible for coordinating access to the PHY, and implementing the basic mechanisms of the interference temperature model. The upper MAC handles higher-level functions like device discovery and authentication. We describe some of the necessary features of the upper MAC, but most of it is left as future work.

4.1 ITMA PHY Layer

The physical layer defines the RF properties of a link between two cognitive radios. It instantiates the bandwidth and power values determined by the interference temperature model.

There are a couple basic underlying communication techniques that would

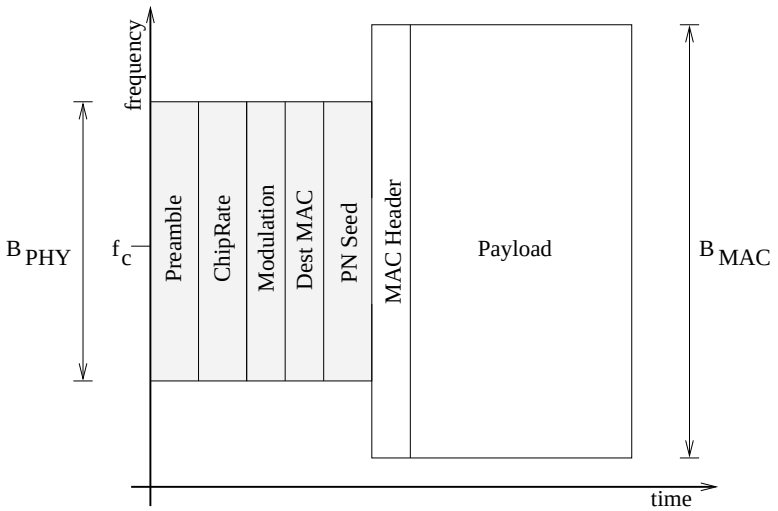

Figure 4.1: Packet spectral occupancy as a function of time, illustrating the dynamic bandwidth used by each packet

lend themselves well to this environment. The first is orthogonal frequency division multiplexing (OFDM). In OFDM, signals are constructed in the frequency domain, rather than the time domain, and therefore creating a signal of arbitrarily large bandwidth can be easily achieved. See Chapter 5 for a more detailed description of OFDM.

Another fundamental underlying waveform, and the one upon which we will focus in this chapter, is direct-sequence spread spectrum (DSSS) [34]. In DSSS the bandwidth of your transmitted signal is generally a function of the chip rate you use to spread your signal. If your chip rate is n times faster than your symbol rate, your spread bandwidth will be n times larger than your narrow-band bandwidth.

As the next section describes, the MAC layer will instruct the PHY on a power and bandwidth to use for a particular transmission that meet the interference temperature model constraints. The PHY must implement the specifications.

Before each packet transmission, the cognitive node determines suitable capacity C and range L requirements for the data to be sent. Typically C will be selected on a per-application basis and L will be selected on a per-destination basis. Using these values, the necessary bandwidth can be computed, which will be accomplished by selecting an appropriate DSSS chip rate.

Each packet is preceded by a PHY header that is spread using one of several well-known pseudonoise (PN) sequence with predefined chip rates. This header contains the PN generator seed and chip rate used for the remainder of the packet. This is illustrated in figure 4.1. For every packet transmitted, the transmitter will compute a new PN sequence. This provides CDMA-like features for the MAC.

Long, non-repeating PN sequences should be used. A good candidate would be simple m-sequences [31]. The space of seeds should be sufficiently large to prevent frequent reuse, which increases the probability that two simultaneously transmitted packets use the same PN sequence and would interfere with each other. A 16-bit value should be sufficient.

On important requirement is that the PHY header not cause harmful interference. If f_c is close to a licensed signal, than using a large bandwidth might cause problems. As a result, for each possible well-known chip rate and spreading code, the associated bandwidth, power, and capacity should be computed. The one maximizing the capacity function should be selected. This offers the highest probability of packet delivery.

The spreading code used for the PHY header should also have a low autocorrelation value. Since there is a possibility of two PHY headers being simultaneously

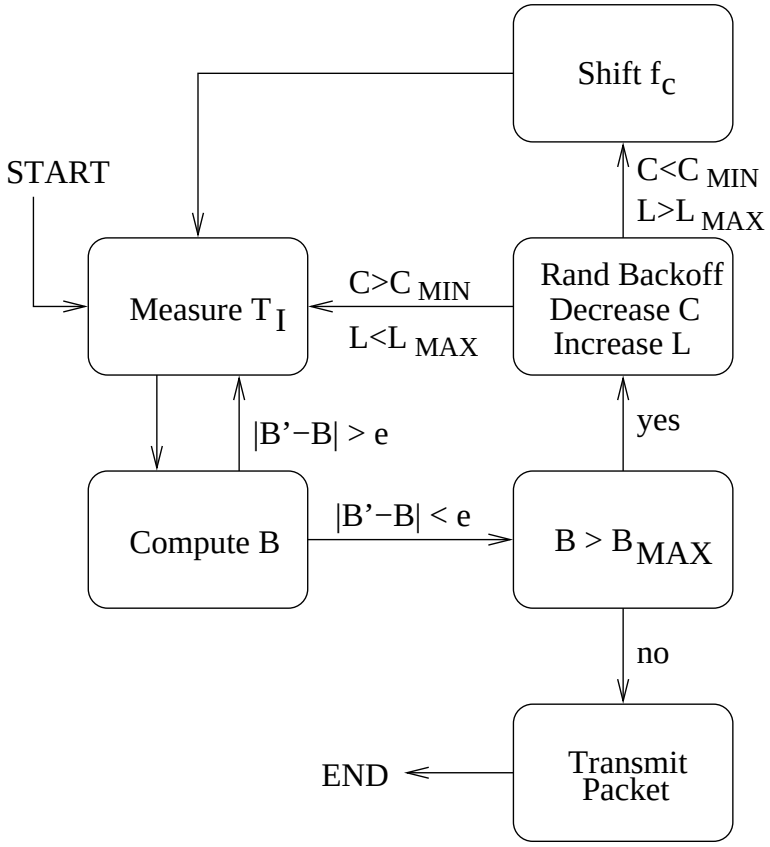

Figure 4.2: General state machine for ITMA, indicating T_I measurement loop with decreasing QoS and eventual frequency shift if unable to transmit packets

transmitted, to minimize interference something like a Barker code should be employed.

Receiving radios must be able to sync up to the preamble and recognize the chip rate and spreading code. This can be done using readily available spread-spectrum technology.

4.2 Basic ITMA MAC Layer

Figure 4.2 depicts the operation of ITMA. When a node wishes to transmit a packet, it first measures the interference temperature, T_I. As described in Chapter

3, this would likely be implemented as an iterative process, since T_I is a function of the measurement bandwidth.

The MAC then uses this value in conjunction with the interference temperature limit T_L, the desired capacity C, and the range parameter L, to compute the required bandwidth B and chip rate. This must be done several times in order as B converges. Given some tolerance ϵ, the cycle repeats until

$$|B_i - B_{i-1}| < \epsilon \tag{4.1}$$

If the bandwidth required to successfully transmit the packet is less than a specified maximum bandwidth $B_{\max}$, the packet is transmitted.

If $B > B_{\max}$, the packet cannot be transmitted. The radio has a few options available, including those that follow.

1. It can simply wait. Transient interference could be causing a temporary inability to communicate. Once T_I decreases, communication can resume.

2. It can decrease C, decreasing the packet's data rate. For services requiring a minimum throughput, it may not be possible to decrease C below a predefined threshold $C_{\min}$.

3. The node can increase L, decreasing the radio's range. If the packet's destination is at a distance or is subject to fading or shadowing, L cannot be increased beyond some maximum $L_{\max}$.

4. If $C < C_{\min}$, $L > L_{\max}$, and some timeout period has expired, the last resort

is to shift the network to a new center frequency. Section 4.3.1 describes this further.

It is important to note the features afforded by using unique spreading codes for every packet transmitted. Initially a receiver need only listen to the PHY header for the packet. If the MAC address contained within the header does not match one of its own, it does not have to demodulate the rest of the transmission. This means that two packets can be transmitted simultaneously without colliding, so long as their headers are disjoint in time. A transceiver can receive a PHY header, and if the address does not match, it can immediately transmit its own packet without waiting for the first packet to finish. The near-far effect can be combated by effective choice of L (see section 4.3.2).

This observation indicates that the system will be more efficient with a large maximum transit unit (MTU). Larger MTUs reduce header overhead. Since headers are the only thing that can collide, the fewer of them transmitted the better.

4.3 Higher MAC Functions

This section describes some of the higher-level MAC functions implemented by ITMA. It addresses techniques for selecting a center frequency, how to better measure T_I and L, and discusses the the hidden terminal problem and how to mitigate its effects.

4.3.1 Center Frequency Selection

While bandwidth and power can change from packet to packet, the center frequency f_c should remain fairly static. When a network is initially configured, an optimal value must be selected. Algorithms for accomplishing this are described in Section 3.3.

If the interference temperature increases to a point at which communication within bandwidth limitations is not possible for the given QoS requirements, a frequency shift may be required. Since significant overhead will be required to regain connectivity among all nodes, frequency shifts should only be used as a last resort.

When network connectivity is lost, a node enters a scan and beacon cycle. In the scan mode, it hops between all $f \in \mathcal{F}$ searching for other nodes with whom it can communicate. At the same time, it records the interference temperature T_I at each frequency. If no other nodes are found, it performs the initial frequency selection and begins beaconing. After some random timeout, if no devices have connected, it resumes the scan mode.

A potential problem is a network partition where multiple radio networks form on different center frequencies. These can be combined in the same way IEEE 802.11 consolidates ad-hoc networks with the same network name.

ITMA intentionally does not include unauthenticated packets that instruct a network to change center frequencies. It would yield a very powerful denial of service attack by flooding spoofed frequency change messages at different locations within

the network. If upper-level security is enabled and devices are authenticated, more coordinated frequency shifts are possible if connectivity isn't completely lost.

4.3.2 Statistics Exchange

Imagine two nodes N_1 and N_2 communicating. In every packet sent from N_1 to N_2, N_1 will include its current T_I and the value of received power P'_S from the last packet from N_2.

These statistics can help N_2 in several ways. First, N_2 can use P'_S to gauge the distance between itself and N_1, which can be used to optimize L and compute $L_{\max}$. If P'_S/kB is significantly higher than T_I, L can likely be increased on packets sent to N_1.

The near-far problem affecting CDMA is a little different with ITMA since bandwidth varies from packet to packet. In CDMA the problem is solved by strict power control. Notice that in ITMA power is fixed, relative to T_I and T_L so we solve it through bandwidth control.

Secondly, T_I can be used to better judge the interference environment at N_1. N_2 could take a weighted average of the surrounding interference environment when computing P_S to avoid causing unintentional interference to spectrum licensees.

4.3.3 Hidden Terminal Problem

For the most part, CDMA has been used in infrastructure networks, typically cellular in nature. Devices only communicate with a base station or access point,

and never to each other. As a result, it would never be the case that two or more devices in the network were simultaneously communicating with a third device.

The hidden terminal problem [32] primarily affects CSMA networks, however moving CDMA to an ad-hoc topology introduces a new problem called the *concurrent transmission problem*. Here, two or more devices in the network can simultaneously transmit to the same destination device. As these transmissions will be appropriately power controlled, the receiver will be capable of receiving any one of the messages. Assuming it has only a single radio and limited DSP power, it cannot receive more than one.

This problem is different from the hidden terminal problem, because in the hidden terminal problem, when a packet collision occurs, all packets are lost. Here, packets are not lost, but only one of the packets can be received. This feature greatly improves ITMA's ability to transmit data; however, while these extra packets do not collide, they raise the overall interference temperature, decreasing our per-link capacity.

To combat this problem, we introduce the idea of *destination hints* to help decrease the chance of a concurrent transmission. In ITMA, radios receive the PHY headers for all packets in their range. These headers contain the address of the destination. Destination hints is an approach where if node A sees a header addressed to node B, it assumes the receiver at node B is busy, and A should not transmit for some amount of time. If we add packet timing information to the PHY header, nodes can know precisely how long a particular receiver will be busy.

This approach would work best in single-hop networks, since everyone in the

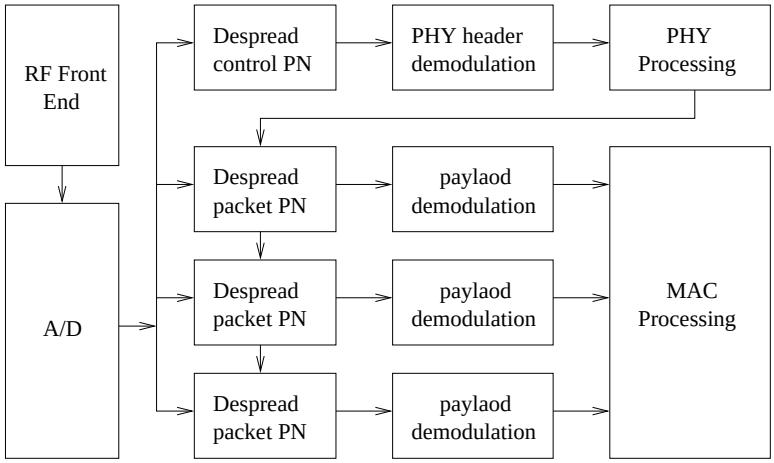

Figure 4.3: Diagram of a radio node capable of receiving three packets simultaneously

network would be able to see everyone else's PHY headers. In multi-hop networks, not all nodes can see all transmissions, so hints would only lessen the problem, not solve it.

Another approach is to have radios that support receiving multiple packets simultaneously, much like a cellular base station. As long as SIR constraints aren't violated, it should be possible to receive all packets. Figure 4.3 diagrams a transceiver that can receive three packets. More than two or three receivers is not likely to help performance, as radios supporting the bandwidth required to properly decode all packets, given equal SIR, would not be cost effective.

4.4 Simple Example

Consider the network depicted in figure 4.4. Nodes are equidistant, spaced 500 meters. Assume we can guarantee there are no licensed receivers within 200 meters of the transceivers. Using free-space path loss, and assuming $f_c = 600$ MHz, the

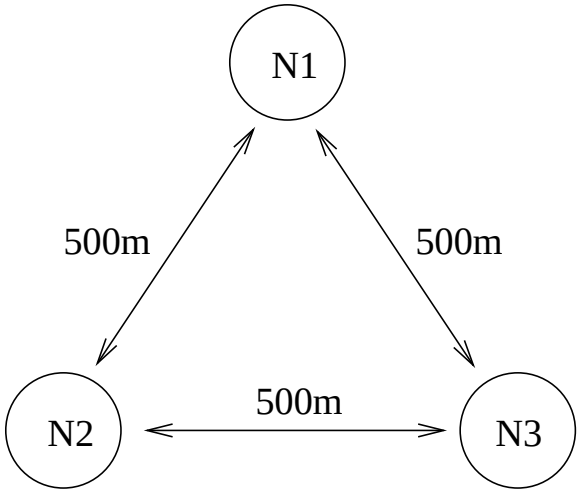

Figure 4.4: Small network of three equidistant nodes

loss variables can be computed as:

$$L = 1 \cdot 10^{-7}$$

$$M = 6 \cdot 10^{-7}$$

For our temperatures, let us use the following values:

$$T_I = \ \ 293 \text{ Kelvin}$$

$$T_L = 3000 \text{ Kelvin}$$

Here, we assume the background interference is caused solely by thermal noise, hence an interference temperature of 293 Kelvin.

Let T_{S_i} be the signal temperature for node i. Each transmitting node will contribute to the base interference temperature T_I, and therefore will cause other nodes to decrease in power. The steady-state powers for all nodes transmitting is

governed by

$$MT_{S_i} = T_L - T_I - L\sum_{j \neq i} T_{S_j} \tag{4.2}$$

Solving this system of equations we obtain the following solution:

$$T_{S_i} = \frac{T_L - T_I}{2L + M} \tag{4.3}$$

To compute the bandwidth required, we need an effective T_I composed of the base interference in addition to the new signals.

$$\tilde{T_I} = T_I + 2L\frac{T_L - T_I}{2L + M} \tag{4.4}$$

Assuming a desired capacity of 5 Mbps, and substituting this into the bandwidth equation we obtain a bandwidth requirement of 11.6 MHz. This equates to a transmit power of -62.7 dBm.

Results on this order of magnitude should be quite acceptable to both spectrum licensees and secondary users. For example, most analog television sets have a receive sensitivity of roughly -50 dBm. These transmissions would therefore never interfere with current broadcast TV. However, in the future, DTV will have much tighter restrictions, with sensitivity on the order of -110 dBm. In this case, a much lower T_L would be necessary to reduce interference.

4.5 Network Analysis

This section presents a mathematical analysis to examine network scalability as a function of node density. For simplicity, we analyze a synchronous network where in each time slice a device can either transmit a packet or not transmit a packet.

Let the probability of transmission be p. Assume the transmission is omni-directional, and can reach a maximum of m neighbors, and it is addressed to one of the neighbors uniformly. Consequently the probability of node A transmitting a packet addressed to particular neighbor node B in any given time slice is p/m.

Next we must consider the effects of collisions. For simplicity, let us assume there is no contention between the PHY headers used in ITMA. A node A successfully receives a packet if at least one of its neighbors transmits a packet addressed to A, and A itself does not transmit. This probability is

$$P_{ITMA}(p, m) = (1 - p)\left(1 - \left(1 - \frac{p}{m}\right)^{m}\right) \tag{4.5}$$

To compute bounds on capacity, we must select a probability $p*$ such that the following holds:

$$P_{ITMA}(p^*, m) \geq P_{ITMA}(p, m) \qquad \forall p \in [0, 1] \tag{4.6}$$

Unfortunately, a closed-form maximizer does not exist. However, we can compute p^* numerically and substitute it to obtain $P^*_{ITMA}(m)$. If we evaluate the limit

numerically, we have

$$\lim_{m\to\infty} P^*_{ITMA}(m) \approx 1/5 \tag{4.7}$$

This results in a constant complexity

$$O\left(P^*_{ITMA}(m)\right) = \mathcal{O}(1) \tag{4.8}$$

As the neighbor count of an ITMA network increases, the probability of a node receiving a packet during each time slice approaches an asymptote. However, we must consider that as the number of neighbors increases, so do the number of concurrent transmissions. As this happens, our SIR decreases, which assuming a fixed maximum bandwidth and power, decreases our per-link capacity. The effects of this are as follows:

$$\begin{aligned} C_{link}(m) &= B\log_2\left(1+\frac{P_S}{mP_I}\right) \\ \mathcal{O}(C_{link}(m)) &= \mathcal{O}(\log(1+1/m)) \\ &= \mathcal{O}(1/m) \end{aligned} \tag{4.9}$$

The last step is due to the fact that a Taylor series expansion of $\log(1+1/m)$ is as follows:

$$\log(1+1/m) = \frac{1}{m} - \frac{1}{2m^2} + \frac{1}{3m^3} - \frac{1}{4m^4} + \cdots \tag{4.10}$$

Application of $\mathcal{O}(\cdot)$ yields $\mathcal{O}(1/m)$. Thus our overall per-node capacity as a function

of its neighbor count is

$$\begin{aligned} C_{ITMA}(m) &= \mathcal{O}(1) \cdot \mathcal{O}(1/m) \\ &= \mathcal{O}(1/m) \end{aligned} \tag{4.11}$$

These results indicate that as node density increases, overall capacity decreases. This should be fairly obvious. As we share a fixed resource among m nodes, the per-node allocation will be $\mathcal{O}(1/m)$. This indicates we want to minimize our neighbor count in order to maximize our per-node capacity. As a network grows, this implies a transition from a single-hop ad-hoc network to a multi-hop mesh-like network.

Consider a multi-hop network of ITMA-based nodes covering a fixed area. As the total number of nodes n increases in this fixed area, so does the average node density.

If we place no restrictions on transmit power and bandwidth, and allow transmission at the radio and regulatory maximums, the number of neighbors will increase with n, that is $\mathcal{O}(m) = \mathcal{O}(n)$. Thus, the per-node capacity is $\mathcal{O}(1/n)$, and the network-wide capacity is

$$\begin{aligned} C_{net}(n) &= \mathcal{O}(1/n) \cdot n \\ &= \mathcal{O}(1) \end{aligned} \tag{4.12}$$

So as our network density increases, if power control is not used, we have a constant overall network capacity as a function of the number of nodes. Also notice

Table 4.1: Summary of network capacity versus network latency trade-off in multi-hop ITMA-based networks, in terms of node count n.

Network	**Capacity**	**Latency**
Minimize Latency	$\mathcal{O}(1)$	$\mathcal{O}(1)$
Hybrid Approach	$\mathcal{O}(\sqrt{n})$	$\mathcal{O}(\sqrt[4]{n})$
Maximize Capacity	$\mathcal{O}(n)$	$\mathcal{O}(\sqrt{n})$

that the network latency remains constant as $\mathcal{O}(1)$, since packets still traverse the same distance with each hop.

To maximize capacity, however, we need to consider a power/bandwidth-controlled scenario that keeps our neighbor count constant as n increases. Thus $\mathcal{O}(m) = 1$, implying a per-node capacity of $\mathcal{O}(1)$. The network-wide capacity is then

$$\begin{aligned} C_{net}(n) &= \mathcal{O}(1) \cdot n \\ &= \mathcal{O}(n) \end{aligned} \tag{4.13}$$

This gives us obvious capacity gains, but we must consider the latency. To maintain a constant neighbor count, the transmission range must decrease as $\mathcal{O}(1/\sqrt{n})$. The result is a latency that increases as $\mathcal{O}(\sqrt{n})$.

Certainly hybrid approaches also exist. One example is where we allow our neighbor count to increase as $\mathcal{O}(\sqrt{n})$. The result is a capacity that grows with $\mathcal{O}(\sqrt{n})$ and a latency of $\mathcal{O}(\sqrt[4]{n})$. These results are summarized in table 4.1. In general, we have

$$\text{capacity} = \mathcal{O}(\text{latency})^2 \tag{4.14}$$

Also, we must note that as node density increases, we will eventually reach a saturation point. Results from the next section show that extremely dense topologies can be easily accommodated.

4.6 ITMA Simulator

In order to evaluate ITMA and some of the various techniques for reducing concurrent transmission, a simulator supporting interference temperature measurements was created. Given the power and SIR seen by a transmission, it computes an information theoretic maximum capacity for each packet. After a certain amount of time has passed, the total number of successful bits is divided by the simulation time to determine a network-wide capacity.

In this simulation, we ignore propagation delay. This is a reasonable assumption as typically propagation delay is small compared to transmission time. A 10 Mbps, 300m link requires 800 μs to transmit a 1000-byte packet, while the propagation delay is 1 μs. This simulator uses the free-space path loss radio propagation model, where power decays as a function of the distance squared, though this is configurable.

4.6.1 MAC Design

The ITMA MAC is broken down into two main events, `tx-start` and `tx-end`. The `tx-start` event is executed whenever a node wishes to transmit a packet. It executes the following tasks:

1. if currently receiving a packet, backoff
2. select a destination node
3. measure the interference temperature
4. compute the power and bandwidth required for the transmission
5. if the required bandwidth is outside the radio or regulatory specifications, backoff
6. set transmission flag
7. if destination node is transmitting, set packet as lost
8. if destination node is receiving as many packets as possible, set packet as lost
9. increment the number of packets the receiver is receiving
10. schedule `tx-end` event

The *backoff* command consists of rescheduling the same `tx-start` event for some randomly chosen time in the future. The `tx-end` event is scheduled for when packet transmission is complete. It executes the following tasks:

1. reset transmission flag
2. decrement the number of packets the receiver is receiving
3. if the packet is not lost, increment the number of received packets
4. schedule a new transmission

At the termination of every scheduled event, the simulator evaluates all current transmissions to see if the interference temperature at each receiver has increased too much.

When a packet is transmitted, B and P_S are computed as a function of some desired capacity αC. Here C is the target capacity, and $\alpha > 1$ is a scaling factor that adds a safety margin. In a real radio, capacity is instantiated by some set of modulation and coding, which cannot be changed in the middle of a packet if interference increases. If we used the minimum B and P_S, a slight increase in the interference temperature at the receiver during packet transmission would prevent reception. Using the scaling factor α protects us from this.

We deem a packet lost if the following inequality does not hold:

$$C \geq B \log_2 \left(1 + \frac{P}{kBT_I}\right) \tag{4.15}$$

where B and P were computed with respect to a desired capacity times a safety margin αC. This is evaluated as follows:

1. loop through all transmitting nodes $t \in T \subseteq N$
2. measure the IT at the receiver
3. compute the capacity C'
4. if $C' < C$, mark packet as lost
5. end loop

4.6.2 ITMA Parameter Experiments

In this section, we simulate ITMA over its parameter sets. This will illustrate how each parameter affects overall network performance and transmit power. All

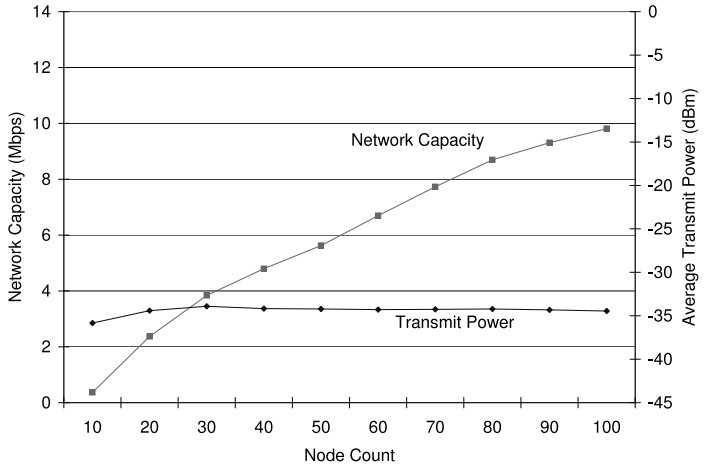

Figure 4.5: Network capacity and transmit power as a function of node density

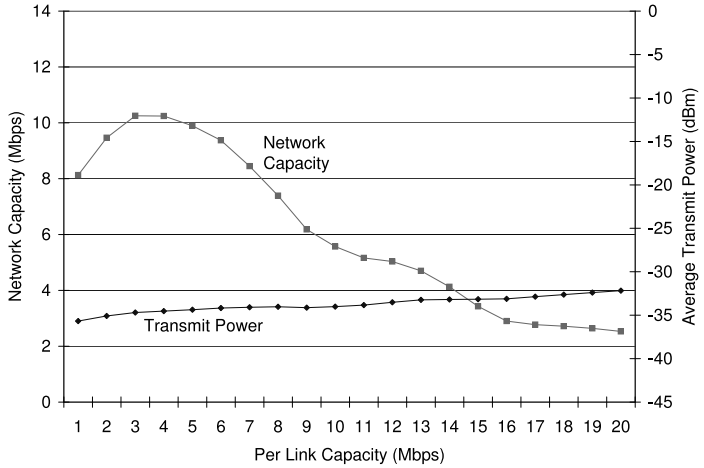

Figure 4.6: Network capacity and average transmit power as a function of the desired per-packet receive capacity

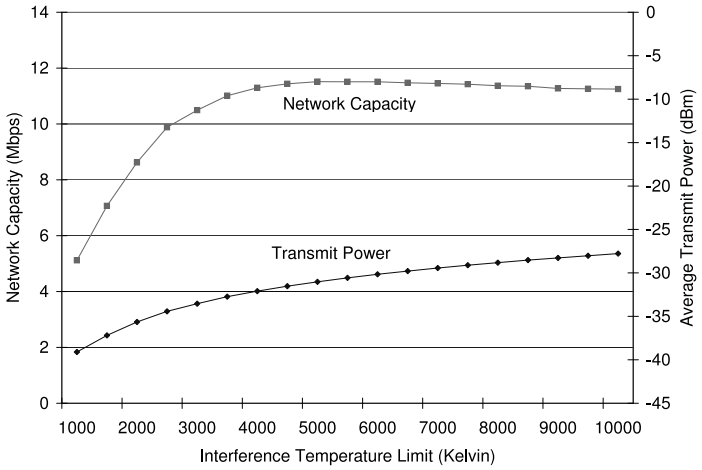

Figure 4.7: Network capacity and average transmit power as a function of the interference temperature limit

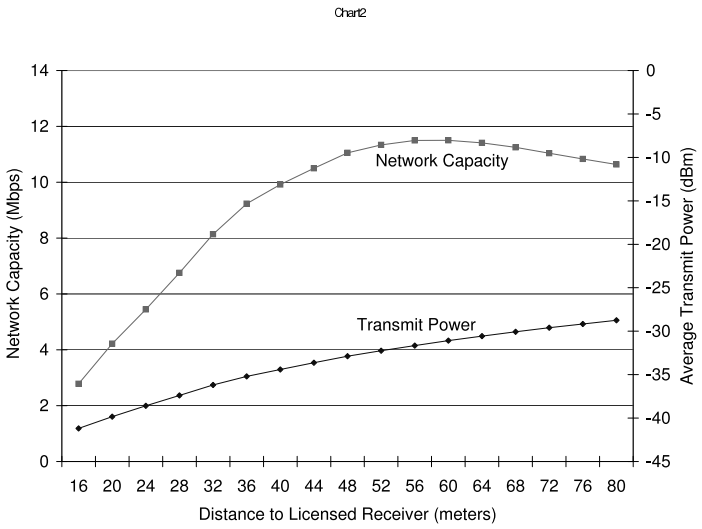

Figure 4.8: Network capacity and average transmit power as a function of M, where M is computed using free-space path loss over the distance specified on the x-axis

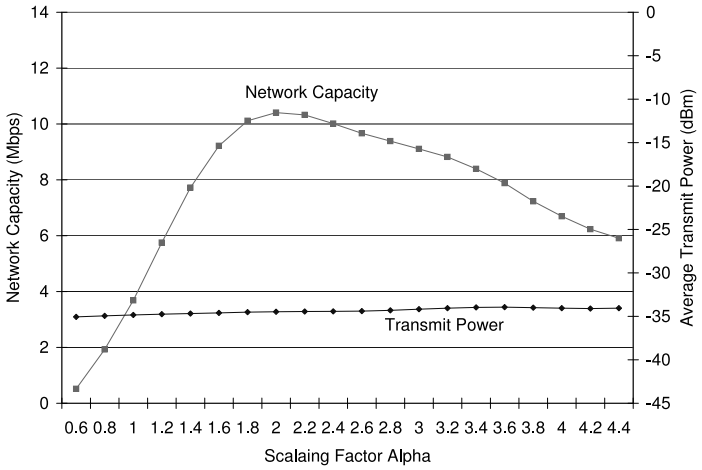

Figure 4.9: Network capacity and average transmit power as a function of the safety scaling parameter α

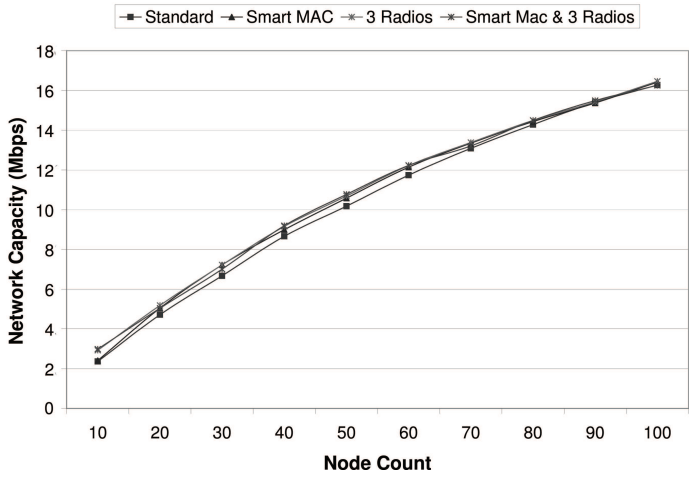

Figure 4.10: Network capacity as a function of node count, plotted with various concurrent transmission mitigation techniques being used

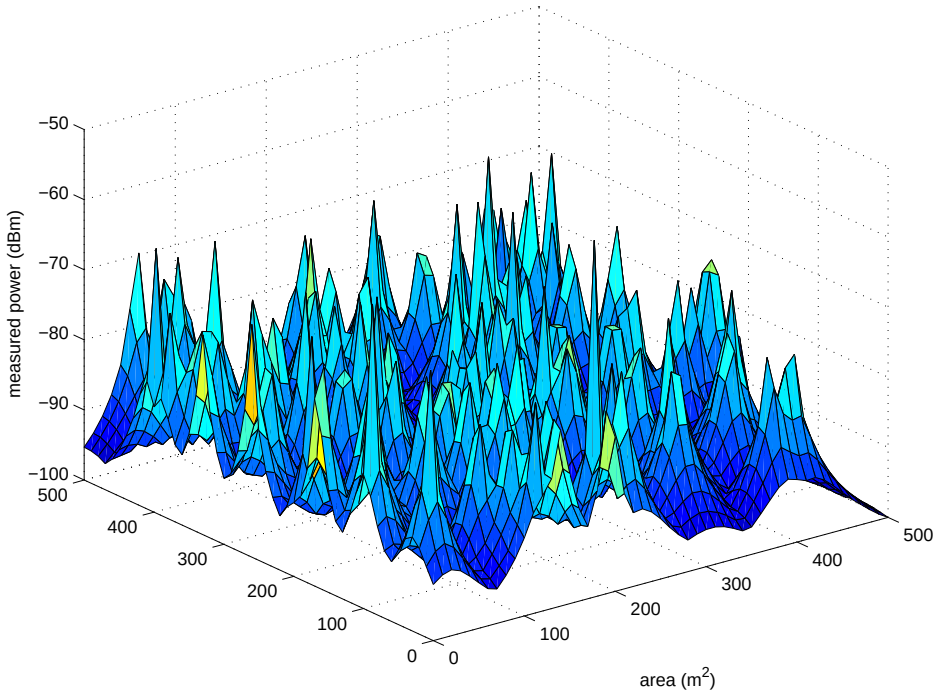

Figure 4.11: Maximum signal powers measured on 50x50 grid over the experiment area.

simulations assume nodes are in a 500 meter by 500 meter world, with positions selected uniformly over the area. We assume a base interference temperature of 293 Kelvin, caused by thermal noise.

To illustrate how various parameters affect performance, we vary each while keeping others constant at a reasonable value. Table 4.2 lists default simulation parameters.

Each plot shows the total network capacity. All plots show both the total network capacity and the average transmit power. The network capacity is computed as the sum of all successfully received bits across the network divided by the simulation time. The average transmit power is the mean power across all packet transmissions.

Table 4.2: Base simulation parameters

Parameter	**Value**
node count n	100
target link capacity C	5 Mbps
IT limit T_L	2500 K
ITMA constant M	40 meters
ITMA constant α	2.5
max radio bandwidth $B_{\max}$	20 MHz
max radio power $P_{\max}$	10 mW
center frequency f_c	600 MHz

Figure 4.5 examines how capacity changes with the increase of the number of nodes spread over the 0.25 km^2 area. This is the network density problem. Impressively, without any special interference mitigation techniques, ITMA achieves near-linear scaling. In fact, regression analysis indicates that the achieved network capacity $C'_{net}(n)$ is

$$\Omega(\sqrt{n}) < C'_{net}(n) < \mathcal{O}(n) \tag{4.16}$$

The simulation results in a hybrid scheme tending toward network capacity optimization. Power is relatively unaffected by node density.

Figure 4.6 shows how capacity and power change with the desired packet receive capacity. Initially, as the desired capacity increases, we have an increase in overall network performance. However, quickly the network is saturated, and we reach a global maximum at roughly 3.5 Mbps. Increased spectrum utilization required to reach the target capacity causes harmful interference to other network users. Interestingly if you look at packet delivery rates, as C increases, they approach 100 percent. Fewer packets are sent, but there are no lost packets.

Figure 4.7 depicts both network capacity and transmit power as a function of T_L. Increasing T_L represents the FCC allowing more interference from unlicensed devices. Keeping T_L small limits the transmit power and consequently the ranges of the radios. The network capacity tops out at just under 12 Mbps, while the average transmit power is $\mathcal{O}(log(T_L))$. It's interesting that capacity doesn't continue increasing with the interference temperature limit increase. If you look at lost packets, as T_L increases, so does the number of transmitted packets. Unfortunately, the number of lost packets increases too, resulting in the asymptotic behavior.

Figure 4.8 reflects network capacity and transmit power as the distance increases between unlicensed transmitters and the licensed devices with which they may interfere. To compute M from the varying distance, the free-space path loss model was used. Increasing the distance reflects an ability to transmit with higher powers. Network capacity maximizes at 56 meters. Transmit power is again logarithmic. Packet loss is similar to increasing T_L.

Lastly, we examine the safety parameter α in figure 4.9. It is used to compute the transmit power and bandwidth necessary for a desired capacity. The higher the value of α, the more bandwidth will be used. This will allow successful packet reception even if the interference temperature increases at the receiver during transmission. Obviously for $\alpha < 1$ successful packet transmission is unlikely, as capacity constraints are almost always violated when the packet is first transmitted. To minimize wasted network resources, a small α should be selected. We see a global maximum at $\alpha = 2.0$.

The major results of these simulations are as follows. First, in almost all cases,

transmit power is logarithmic in each of the parameters. Also, for two parameters, C and α, we can find unique, global maximizers. Both parameters are local to each radio, so cognitive nodes could update them in real time to optimize overall network performance. Lastly, these simulations provide insight into good ways to select the FCC-controlled parameters M and T_L.

4.6.3 Concurrent Transmission Mitigation Experiments

In section 4.3.3 we introduced some techniques for mitigating the effects of concurrent transmissions. The first is a *Smart MAC* that implements both destination hints, and also senses whether or not the destination is transmitting. The second is to have a radio node with multiple receivers.

To investigate the impact these techniques have on a network of ITMA nodes, we included support for these in our simulator. Figure 4.10 plots network capacity for the various techniques as a function of network density. We can see that for sparse networks, we can achieve a 30% performance increase by implementing these techniques. However, as the node count increases, the advantage decreases. At 100 nodes, there is only a 1% increase in overall performance when the concurrent transmission mitigation techniques are used.

It is expected that deterministic networks with relatively static traffic patterns would benefit more from these mitigation techniques.

4.6.4 Interference Analysis

In this section we investigate the amount of interference caused by an ITMA-based network, and how it relates to the interference temperature limit. To accomplish this, a measurement routine was added to the simulator. It is responsible for measuring the signal power on a grid with 10 meter edges, throughout the experiment area. At the end of the simulation, it outputs the maximum power recorded at each measurement site.

Figure 4.11 shows these signal powers in a three-dimensional plot, when the simulator is executed using the parameters from table 4.2. Spikes represent the locations of cognitive radio transceivers. The measured powers range from -100 dBm to -60 dBm.

This chapter provided an analysis of the scalability and capacity that can be achieved by an ITMA-based network. It shows that ITMA scales such that $\mathcal{O}(\text{capacity}) = \mathcal{O}(\text{latency})^2$. It also examines some techniques for mitigating the concurrent transmission problem, and evaluates their performance through simulation.

Chapter 5

Spectral Shaping

The fundamental concept of the interference temperature model is to avoid raising the *average* interference power for some frequency range over some limit. However, if either the current interference environment or the transmitted underlay signal is particularly nonuniform, the *maximum* interference power could be particularly high.

Another interesting application of spectrum shaping would be in the ideal interference temperature model, where we could regulate control interference power when we overlap licensed signals, but transmit with higher power at other frequencies.

5.1 Problem Formulation

To better understand the problem, let's examine the differences between maximum power and average power. Figure 5.1 depicts three power curves, interference $P_I(f)$, signal $P_S(f)$, and total $P_T(f)$, where

$$P_T(f) = P_I(f) + P_S(f) \qquad \forall\ f \in [f_c - B/2, f_c + B/2] \tag{5.1}$$

Each signal has mean over bandwidth B of $\bar{P}_I$, $\bar{P}_S$, and $\bar{P}_T$, respectively.

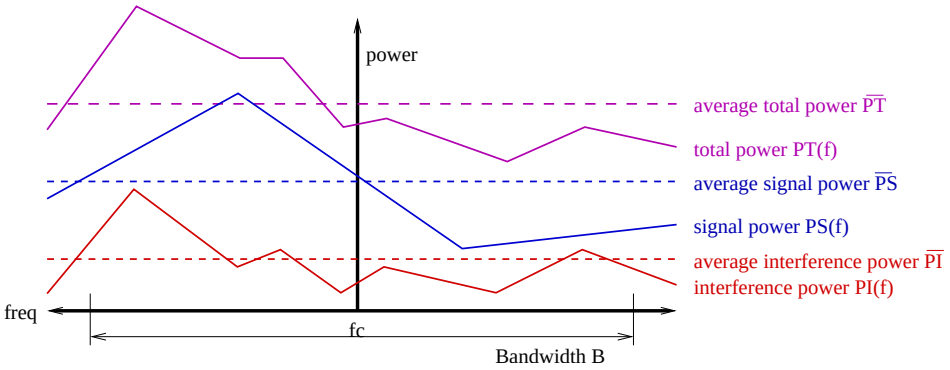

Figure 5.1: Figure showing that approximating $P_T(f)$ over the interval $[f_c - B/2, f_c + B/2]$ by the average power $\bar{P}_T$ could yield unexpected interference exceeding regulatory allowances..

Assuming a fixed bandwidth, then the interference temperature model stipulates that

$$\bar{P}_T \leq B\,k\,T_L \tag{5.2}$$

As figure 5.1 indicates, even with equality we have no real guarantee on absolute maximum interference. A stronger requirement would be

$$P_T(f) \leq B\,k\,T_L \qquad \forall\ f \in [f_c - B/2, f_c + B/2] \tag{5.3}$$

Note that this requirement wholly implies the first. That is,

$$\max P_T(f) \geq \bar{P}_T \quad \Rightarrow \quad \big(P_T(f) \leq B\,k\,T_L \quad \Rightarrow \quad \bar{P}_T \leq B\,k\,T_L\big) \tag{5.4}$$

Thus, in order to maximize both capacity and spectral efficiency while minimizing absolute interference, we must find some $P_S(f)$ such that $P_T(f) = B\,k\,T_L$.

We can actually build off the algorithms developed in Chapters 3 and 4. These

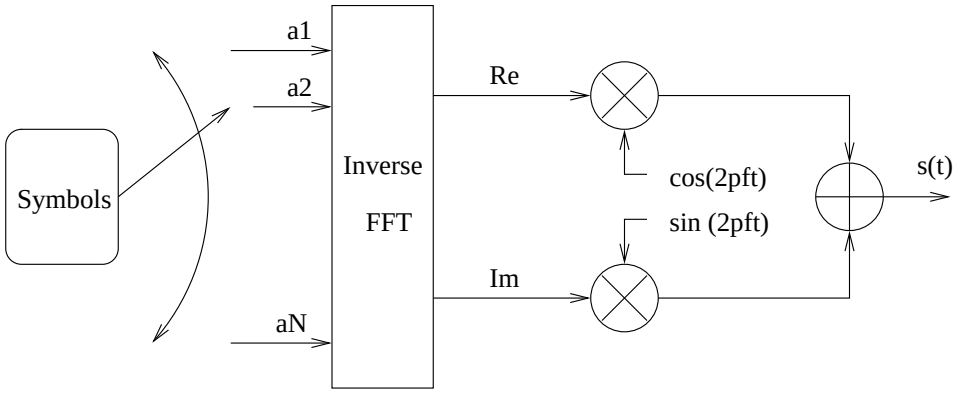

Figure 5.2: Simplified OFDM transmitter.

algorithms give us $\bar{P}_S$ and B[1] that reach our target capacity while satisfying the interference temperature constraints. Knowledge of B is sufficient, and we compute

$$P_S(f) = B\,k\,T_L - P_I(f) \qquad \forall\ f \in [f_c - B/2, f_c + B/2] \tag{5.5}$$

Now that we know the bandwidth and spectral shape required of our transmission, the goal is to actually create a waveform that achieves it. The next sections describe two ways of accomplishing our goal. The first technique uses power control across OFDM sub-carriers. The second approach creates spreading codes with certain spectral characteristics that shape the signal.

5.2 Spectral Shaping with OFDM

In OFDM, waveforms are essentially constructed in the frequency domain and converted into the time domain before being transmitted. Figure 5.2 illustrates such

[1]Note that B either reflects B_L or B_U depending on whether we are considering the ideal or generalized model, respectively.

a transmitter. Symbols are complex numbers representing modulated bit streams. For example, if quadrature phase shift keying (QPSK) is the underlying modulation technique, bit streams of length two would be modulated as $1+j$, $1-j$, $-1+j$, or $-1-j$. The complete block $\{a_1, ..., a_N\}$ are then sent into the inverse Fast Fourier Transform (FFT) to produce a time-domain waveform. The real and imaginary components of the complex baseband signal are multiplied by sin and cos to create the passband signal $s(t)$. Mathematically, the complex-baseband signal $v(t)$ can be expressed as

$$v(t) = \sum_{k=1}^{N} a_k \mathbf{e}^{j2\pi kt/T} \qquad \forall\ 0 \le t \le T \tag{5.6}$$

Assuming a uniform distribution over the input symbols, the average power spectrum is essentially flat as a function of frequency. Our goal is to affect this average power. In particular, assume our bandwidth B is broken up into N subcarriers, as described. The desired average power for subcarrier k is

$$p_k = B\, k\, T_N - i_k \tag{5.7}$$

where

$$i_k = \frac{N}{B} \int_{f_c - B(N-2k)/2N}^{f_c - B(N-2k-2)/2N} P_I(f)\, df \tag{5.8}$$

Let α be the average symbol power for the underlying modulation scheme. We can then reformulate our complex-baseband OFDM signal as

$$v(t) = \frac{1}{\alpha\sqrt{2}} \sum_{k=1}^{N} a_k p_k \mathbf{e}^{j2\pi kt/T} \qquad \forall\ 0 \le t \le T \tag{5.9}$$

This will convert the average power on subcarrier k from α to p_k. The $\sqrt{2}$ is necessary to normalize since multiplication by p_k will affect both the real and complex portions of the waveform.

Alternatively, symbols could all be multiplied by some relative scaling value such as $p_k/(B\,k\,T_N - \bar{P}_I)$ and then the final signal $s(t)$ could be adjusted such that its average power was $\bar{P}_S$. This approach would likely make more sense in a real-world transmitter where amplification happens in the RF front end.

Regardless, we can now shape our power spectrum. However, we must be able to effectively utilize our spectral resources if we hope to achieve channel capacity. In particular, the capacity on each subcarrier varies, as each has a different SIR. This means different coding is necessary on each subcarrier to achieve capacity.

The capacity on subcarrier k is

$$C_k = B/N \log_2(B\,k\,T_L/i_k) \tag{5.10}$$

This espression is based on the Shannon-Hartley theorem. Its vailidity is based on the assumption that the subcarrier bandwidth is much smaller than the coherence bandwidth for the interference, and as a result we can assume the interference is white.

The total capacity is

$$\begin{aligned} C &= \sum_{k=1}^{N} B/N \log_2(B\,k\,T_L/i_k) \\ &= B \log_2(B\,k\,T_L) - \frac{B}{N}\sum_{k=1}^{N} \log_2 i_k \end{aligned} \tag{5.11}$$

If we have uniform interference, i.e. $i_k = \bar{P}_I$, then this equals the capacities derived earlier for the interference temperature model. However, variances in i_k will actually help us achieve higher capacities, since

$$\frac{1}{N}\sum_{k=1}^{N} \log_2 i_k \leq \log_2 \bar{P}_I \tag{5.12}$$

Thus, with proper channel coding, we can outperform the standard interference temperature model by performing spectral shaping with OFDM. Not only can we decrease the maximum interference experienced by others, but we can also increase our capacity while using the same average transmission power and meeting regulatory requirements.

5.3 Spectral Shaping with DSSS

Direct Sequence Spread Spectrum (DSSS) is another technique commonly employed for creating wide-band signals. Here we start with a narrow-band complex-baseband signal and "spread" it using a spreading code.

In Figure 5.3 we can see the basic operation of a DSSS transmitter. Modulated symbols $a(t)$ are multiplied by a high-frequency signal $c(t)$ before being up-converted

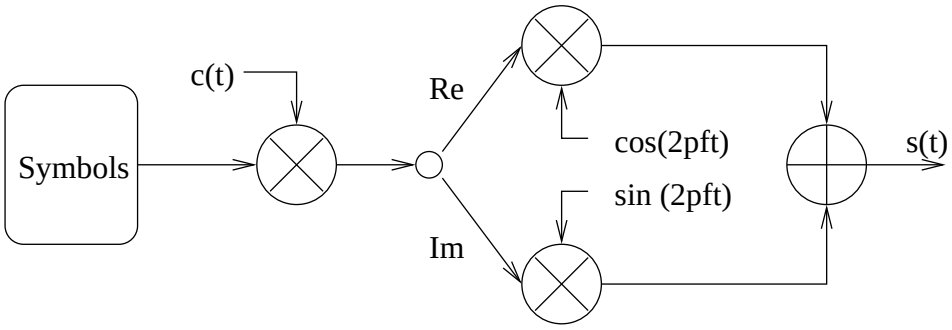

Figure 5.3: Simplified DSSS transmitter.

to passband.

Our goal is to specify $c(t)$ such that $s(t)$ has the spectral properties we desire. Mathematically, we have can define our complex baseband signal $v(t)$ as

$$v(t) = c(t)\, a(t) \tag{5.13}$$

Thus we have

$$c(t) = v(t)/a(t) \tag{5.14}$$

Let's look at $v(t)$, the desired signal. In the frequency domain, we have

$$\begin{aligned} V(f) &= P_S(f - f_c) \\ &= \Pi_B(f - f_c)(B\,k\,T_L - P_I(f - f_c)) \\ &= \Pi_B(f - f_c)(B\,k\,T_L - \mathcal{F}(i(t))) \end{aligned} \tag{5.15}$$

Here Π_B is the width-B rectangular function, and $i(t)$ is the current interference environment, downsampled to passband.

Returning to the time domain, we have

$$\begin{aligned} v(t) &= \mathcal{F}^{-1}(V(f)) \\ &= \mathcal{F}^{-1}\left(\Pi_B(f - f_c)(B\,k\,T_L - \mathcal{F}(i(t)))\right) \\ &= w(t) \star (B\,k\,T_L\,\delta(t) - i(t)) \\ &= w(t) \star (B\,k\,T_L\,\delta(t)) - w(t) \star i(t) \\ &= w(t)\,B\,k\,T_L - w(t) \star i(t) \end{aligned} \tag{5.16}$$

where $\star$ is convolution and

$$w(t) = \frac{\sin(Bt/2)}{\pi t}\,\mathbf{e}^{jf_ct} \tag{5.17}$$

Thus combining everything, we have

$$c(t) = \frac{w(t)\,B\,k\,T_L - w(t) \star i(t)}{a(t)} \tag{5.18}$$

This approach has some major realization drawbacks. Notice that $H(c(t)) \geq H(a(t))$, thus our spreading sequence actually contains *more* information than our information sequence. When we multiply $c(t)$ by $a(t)$, we actually cancel out the data symbols and transmit a signal with exactly the spectral characteristics we want. No actual data flows over the main channel, and everything passes through the side channel in which we convey the spreading code. Thus, this *ideal* approach is not realistic.

As a result, we must assume $i(t)$ and $a(t)$ are stationary, and sample them over

a short period of time. From them, we compute $c(t)$ and discretize it into something we can represent in a finite number of bits to be communicated via our side channel. Such sampling will obviously degrade our performance, but is necessary to make the scheme practical.

Let τ_c be our spreading code's chip time and τ_s be our symbol time. If B_N is our narrow-band bandwidth, we must have

$$\tau_c > B_N \, \tau_s / B \tag{5.19}$$

in order to provide enough bandwidth expansion.

Thus, we must sample $c(t)$ every τ_c units of time, and we need at least τ_s/τ_c samples. More samples will provide a more accurate estimate and decrease interference. Any fewer and we won't get the necessary bandwidth expansion.

Assume we sample both the real and complex values of $c(t)$ with M-bit resolution. Our entire spreading code can be represented in a minimum of $MB/4B_N$ bytes. While this is not insignificant, it could be easily conveyed by a side channel, or the ITMA PHY header. For example, with 16 bytes of data we could accommodate 4-bit quantization for spreading a 2 MHz narrow-band signal to a 32 MHz wide-band signal.

Unlike OFDM, our capacity will remain unchanged. Each symbol is multiplied by a spreading code which may amplify some portions of the symbol and attenuate others. However, the average symbol power will remain unchanged, as compared to a traditional spreading code.

5.4 Conclusion

Overall, these spectrum shaping techniques can help us "fill the regulatory gaps" in a particular interference environment. While the proposed FCC regulations only stipulate average interference over the transmission bandwidth, we can actually achieve the same or greater capacity by shaping our spectra.

This chapter presented an initial, motivational analysis of applying OFDM and DSSS to ITMA. Much research still needs to be done on simulating and testing these ideas. In particular, a complete analysis of how quantization of our spreading codes affects the eventual waveform will yield important results on the viability of this approach.

Chapter 6

Conclusion

This dissertation has examined the problem of dynamic spectrum access in the presence of a licensed signal, when unlicensed communicating devices have intelligent radios capable of sensing and reacting to their environment. In this chapter we review the major results and discuss areas of further research.

The interference temperature model, as proposed by the FCC, could offer a new paradigm for spectrum access, though many telecommunications companies have been reluctant to jump on-board until a full analysis of possible interference effects can be performed. Research presented here should provide a foundation for such analysis. In particular, we can quantify how interference temperature limits should be selected, and how those choices affect the range of licensed signals. Here we examined a simplex, omnidirectional transmission, but these results could just as easily be extended to duplex, directional transmitters. The mathematics may be cumbersome, but using the results presented in this dissertation, simulators could be easily constructed.

In particular, it has been shown that measuring interference temperature is a tricky task. Proved are techniques that can compute a precise transmit power and bandwidth that meet a target capacity while also satisfying the requirements of the interference temperature model. This could be extended by examining the pricing

schemes that trade off the utility of capacity with the cost of bandwidth. Since capacity is a nonlinear function of bandwidth, the resulting pricing schemes could have very interesting results.

Additionally an analytical evaluation of the interference temperature model was completed for both WAN and LAN wireless mesh networks. For both, the capacity achieved is a simple function of the number of nodes, the average bandwidth, and the fractional impact to the licensed signal's coverage area, and scales as $\mathcal{O}(n)$, where n is the number of nodes in the network.

Next, Interference Temperature Multiple Access is introduced. ITMA is a PHY and MAC protocol suitable for implementation on a cognitive radio that supports the interference temperature model. It works by first sensing the RF environment and determining what bandwidth and power are necessary to communicate with a desired capacity. If these parameters are not supported by the radio it either lowers its capacity expectations or searches for a new center frequency that has fewer interference problems. Packets are preceded by a PHY header that is transmitted with known modulation parameters. This header contains the information necessary to demodulate the rest of the packet.

A simulator for ITMA was implemented to test the MAC scheme in a mesh topology. Results show that very realistic WAN and LAN-type applications can easily be supported by ITMA while using transmit powers on the order of -35 dBm. The many parameter trade-offs are examined. Future work could involve a more detailed simulator for ITMA that would allow the development and refinement of its upper MAC-level protocols. This could be done by implementing ITMA in a

simulator such as ns-2.

The interference temperature model deals with *average* powers. However, placing bounds on average power does not necessarily constrain *maximum* power. To mitigate problems caused by transient interference spikes, spectral shaping techniques were also investigated. The basic idea to shape our average power spectrum to precisely complement that of the interference environment, allowing us to not only raise the interference temperature precisely to the interference temperature limit, but also raise the absolute interference power spectrum as close to the limit as possible.

Implementing this idea for both OFDM and DSSS was discussed. For OFDM we use power control on each OFDM sub-carrier. The result is an actual increase in capacity since we can more closely target our channel coding to fit the SIR of each sub-carrier. For DSSS we generate spreading codes with particular frequency-domain characteristics that will achieve the desired signal spectrum. The main difficulty is determining how to quantize the spreading code so it can be included in the ITMA PHY header. Much future work exists in this area. In particular, we are only considering the power spectrum at the receiver, not transmitter. Since the transmitter and interferer are geographically separate, the interference power spectrum at the receiver could be much different than anticipated.

Another area for future work is to examine the effects of fading on the interference temperature model. Initial deployments will likely be stationary, and therefore subject only to multipath fading. Modeling and simulating this can be accomplished by simply increasing the path loss constant. Mobility will offer a more challenging

analysis environment. Here, devices will be subject to the Doppler fading, which results in multiplicative interference. This results in both decreased SIR, and less reliable interference temperature measurements. For these environments, distributed, cooperative spectrum sensing may be required.

Before this research began, the interference temperature model had been proposed, but nobody was sure exactly how it might work. The research presented in this dissertation has filled the gaps in the model itself, and positively analyzed its viability. ITMA is the first concrete protocol ever proposed to use the interference temperature model, and this research has proved that it will work. While spectrum shaping techniques have been previously researched, this particular application represents novel work in the field. This dissertation and the research contained within has significantly advanced the state of the art in cognitive radio communications and in particular the interference temperature model.

Appendix A

Abbreviations

API	application programming interface
AWGN	additive, white, Gaussian noise
BPSK	binary phase-shift keying
CDMA	code-division multiple access
CR	cognitive radio
CSMA	carrier-sense multiple access
DSP	digital signal processor
DSSS	direct-sequence spread spectrum
DTV	digital television
FCC	Federal Communications Commission
FDMA	frequency-division multiple access
FFT	fast Fourier transform
FM	frequency modulation
FPGA	field-programmable gate array
HDL	hardware description language
ISM	industrial, scientific, and medical
IT	interference temperature
ITM	interference temperature model
ITMA	interference temperature multiple access

LAN	local area network
MAC	medium access control
MTU	maximum transission unit
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PC	personal computer
PDA	personal digital assistant
PHY	physical
PN	pseudonoise
PSD	power spectral density
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying
RAN	radio access network
RF	radio frequency
SDR	software defined radio
SIR	signal to interference ratio
TDMA	time-division multiple access
TV	television
UHF	ultra-high frequency
VHF	very-high frequency
WAN	wide area network
WLAN	wireless local area network
WWAN	wireless wide area network

BIBLIOGRAPHY

[1] Agee, B., Kleinman, R., and Reed, J. Soft synchronization of direct-sequence spread-spectrum signals. *IEEE Transactions on Communications* (1996).

[2] Atkinson, K. *Numerical Analysis*. John Wiley and Sons, 1989.

[3] Bertsekas, D., and Gallager, R. *Data Networks*. Prentice Hall, 1992.

[4] Brodersen, R., Wolisz, A., Cabric, D., Mishra, S., and Willkomm, D. Corvus: A cognitive radio approach for usage of virtual unlicensed spectrum. UC Berkeley White Paper, 2004.

[5] Cabric, D., Mishra, S., Willkomm, D., Brodersen, R., and Wolisz, A. A cognitive radio approach for usage of virtual unlicensed spectrum. IST Mobile Summit 2005.

[6] Challapali, K., Mangold, S., and Zhong, Z. Spectrum agile radio: Detecting spectrum opportunities. IEEE GlobeCOM 2004.

[7] Clancy, T., and Arbaugh, W. Measuring interference temperature. Virginia Tech Symposium on Wireless Personal Communications 2006.

[8] Clancy, T., and Arbaugh, W. Stochastic analysis of the interference temperature model. under submission, IEEE ISIT 2006.

[9] Cover, T. M., and Thomas, J. A. *Elements of Information Theory*. John Wiley and Sons, 1991.

[10] Federal Communications Commission. Establishment of interference temperature metric to quantify and manage interference and to expand available unlicensed operation in certain fixed mobile and satellite frequency bands. ET Docket 03-289, Notice of Inquiry and Proposed Rulemaking.

[11] Federal Communications Commission. Facilitating opportunities for flexible, efficient, and reliable spectrum use employing cognitive radio technologies. ET Docket 03-108, Notice of Inquiry and Proposed Rulemaking and Order.

[12] French, C., and Gardner, W. Spread spectrum despreading without the code. *IEEE Transactions on Communications* (1986).

[13] IEEE 802.22. Working group on wireless regional area networks. http://www.ieee802.org/22/.

[14] Ileri, O., Samardzija, D., Sizer, T., and Mandayam, N. Demand responsive pricing and competitive spectrum allocation via a spectrum server. IEEE DySPAN 2005.

[15] Leaves, P., Ghaheri-Niri, S., Tafazolli, R., Christodoulides, L., Sammur, T., Stahl, W., and Huschke, J. Dynamic spectrum allocation in a multi-radio environment: Concept and algorithm. Conference on 3G Mobile Communication Technologies 2001.

[16] Ma, L., Han, X., and Shen, C. Dynamic open spectrum sharing MAC protocol for wireless ad hoc networks. IEEE DySPAN 2005.

[17] Mitola, J. The software radio architecture. *IEEE Communications Magazine* (1995).

[18] Mitola, J. *Software Radio Architecture*. John Wiley and Sons, 2000.

[19] Mitola, J. Cognitive radio: An integrated agent architecture for software defined radio. Ph.D. Dissertation, KTH, 2000.

[20] Mitola, J., and Maguire, G. Cognitive radio: Making software radios more personal. *IEEE Personal Communications* (1999).

[21] Nie, N., and Comaniciu, C. Adaptive channel allocation spectrum etiquette for cognitive radio networks. IEEE DySPAN 2005.

[22] OverDRiVE Website. http://www.ist-overdrive.org/.

[23] Pursley, M. B. The role of spread spectrum in packet radio networks. *Proceedings of the IEEE* (1987).

[24] Raman, C., Yates, R., and Mandayam, N. Scheduling variable rate links via a spectrum server. IEEE DySPAN 2005.

[25] Reed, J. H. *Software Radio: A Modern Approach to Radio Engineering*. Prentice Hall PTR, 2002.

[26] Russell, S., and Norvig, P. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2002.

[27] Satapathay, D., and Peha, J. Etiquitte modification for unlicensed spectrum: Approach and impact. IEEE VTC 1998.

[28] Satapathay, D., and Peha, J. Performance of unlicensed devices with a spectrum etiquitte. IEEE GlobeCom 1997.

[29] Satapathay, D., and Peha, J. Spectrum sharing without licenses: Opportunities and dangers. TPRC 1996.

[30] Shannon, C. Communication in the presense of noise. *Proceedings of the Institute for Radio Engineers* (1949).

[31] Simon, M. K., Omura, J. K., Scholtz, R. A., and Levitt, B. K. *Spread Spectrum Communications Handbook*. McGraw-Hill Professional, 2001.

[32] Tobagi, F., and Kleinrock, L. Packet switching in radio channels: Part II—the hidden terminal problem in carier sense multiple access and the busy tone solution. *IEEE Transactions on Communications* (1975).

[33] Toenjes, R., Benko, P., Ebenhard, J., Frank, M., Goransson, T., Hasmann, W., Huschke, J., Lomar, T., Paila, T., Sallstrom, F., Stahl, W., and Xu, L. Architecture for a future generation multi-access wireless system with dynamic spectrum allocation. IST Mobile Summit 2000.

[34] Viterbi, A. *CDMA: Principles of Spread Spectrum Communication*. Addison Wesley Longman Publishing Co, 1995.

[35] Zheng, H., Zhao, J., Zhang, Z., and Zhao, B. Distributed coordination in dynamic spectrum allocation networks. IEEE DySPAN 2005.