LEABHARLANN CHOLÁISTE NA TRÍONÓIDE, BAILE ÁTHA CLIATH
Ollscoil Átha Cliath

TRINITY COLLEGE LIBRARY DUBLIN
The University of Dublin

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

Copyright statement

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

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

Liability statement

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

Access Agreement

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

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

A Generalized Method for Quantification of Cognitive Radio Benefits within Wireless Systems

A thesis submitted for the degree of
Doctor of Philosophy in Electronic & Electrical Engineering
University of Dublin, Trinity College
Department of Electronic & Electrical Engineering
October 9th, 2009

Preston Fairfax Marshall

TRINITY COLLEGE
1 0 APR 2012
LIBRARY DUBLIN

9383

Declaration

I hereby declare that:

(a) This thesis has not been submitted as an exercise for a degree at this or any other University,

(b) This thesis is entirely the work of the author, except where otherwise cited,

(c) The Trinity College Library may loan and copy this thesis upon request.

SUMMARY

For most of the history of wireless communications, the primary enabler of enhanced radio communications capability and utility has been continual improvement in the implementing technology and the corresponding ability to exploit a wider range of frequencies, provide complex signaling, and perform coding and error correction.

The advent of cognitive radio has introduced an opportunity to further extend the performance of wireless systems. However, most descriptions of cognitive radio have stressed the qualitative benefits of the technology, and have provided a behavioral foundation for the techniques that might be employed. This does not provide a quantitative basis by which their inclusion can be evaluated and justified, in comparison to more conventional approaches for achieving design objectives.

This dissertation develops closed-form expressions of both low and high energy spectrum distributions in order to quantify the ability of cognitive radio to address critical issues in wireless performance and reliability. Expressions of overload probability and intermodulation noise level are derived in order to compare cognitive and non-cognitive radio performance, and to determine the opportunity to reduce required component performance. Similar analysis is conducted on the supportable density of cognitive and non-cognitive radio networks, quantifying the potential increase in density that dynamic adaptation can provide.

The contribution of Dynamic Spectrum Access (DSA) is of particular importance to this analysis. DSA, as a technique to share spectrum without causing interference to other users, has become an important research topic. This dissertation applies DSA with the objective of maximizing the benefit to the user, owner, and operator of a cognitive radio and network through adaptation to the environment, as well as to other spectrum users. It shows that maximal capacity of aggregations of cognitive radios is achieved in conditions in which significant levels of mutual interference are present, and must be tolerated and mitigated.

Six overall conclusions arise from the work reported in this dissertation.

1. The environments of a cognitive radio can be described by a set of closed-form probability distributions that are directly derivable from spectrum measurements, and closely approximate the measured environmental characteristics.

Environments can also be synthesized based on interpolation between, or extrapolation from, measured environments.

2. These closed-form environmental expressions lead directly to estimates of the environmental sensitivity of certain aspects of cognitive and non-cognitive radio performance by providing probability distributions of environmentally sensitive performance metrics.

3. A relatively straightforward set of environmentally informed decisions can significantly improve the performance of wireless devices in a DSA regime. This performance benefit can be realized through a mix of benefits: greatly increased probability of successful operation, lower cost and energy consumption wireless devices due to reduction of hardware performance levels, and through a mix of both performance improvements and cost and energy reduction.

4. The use of DSA changes more than spectrum management. By enabling devices to adapt to their environment, it radically changes the optimal approach to maximize aggregate spectrum utility. Active spectrum re-use maximization requires fundamentally different strategies than those appropriate to minimize the spectrum usage of individual communications links.

5. While non-interfering spectrum sharing with non-cooperative devices is a worthy objective for DSA, significant additional improvements in performance and density are possible when the transition is made to spectrum sharing regimes in which all users of the spectrum can assume that the other devices are capable of interference tolerant operation.

6. Cognitive radio, with DSA functionality, is also enabling of significant adaptation in the upper layers of the network; which can become a source of future performance benefits.

Acknowledgements

First of all, I wish to thank my Doctoral advisor, Linda Doyle, who steered me to both an accomplishable task, and to completion through intercontinental review sessions, meetings, and telephone calls. Equally, I thank my first graduate advisor, Don Hillman of Lehigh University for his preparation in formal methods and information science, and for encouragement during the development of this work.

The bulk of the research reported here was performed while I was at the Defense Advanced Research Projects Agency (DARPA). I thank the DARPA leadership for encouraging this effort, including Tony Tether, Bob Leheny, and Dave Honey.

Cognitive radio is about new ways to use new technology to solve what are sometimes old problems. I appreciate Ulrich Rhode, who brought me an understanding of the ultimate limitations in circuit-based solutions to front-end linearity; Gabriel Rabeiz, whose research into Microelectromechanical Systems filters promises future low-cost, small volume, and high Q filters that instigated my original contemplation of radios utilizing them in lieu of front-end linearity; and Vint Cerf, whose work on Delay Tolerant Networks fits so well into the context of adaptive wireless systems.

I wish to thank Mark McHenry, whose DARPA XG program research will enable the Dynamic Spectrum Access regimes that this dissertation assumes; and who provided me with a hard-drive of raw data from his National Science Foundation sponsored spectrum measurements. These are the basis for quantifying the environments provided in this work.

I wish to extend sincere thanks to many members of the cognitive radio community that were supportive of this effort: Bruce Fette, who provided extensive review and commentary in the context of editing some of this work into several book chapters; Jeff Reed, who always had thoughts and observations to offer in the many meetings we have had over the years; Joe Evans and Gary Minden, who bridged the hardware and software gap; and the many members of the Dynamic Spectrum Access Networks (DYSPAN) community, including Paul Kolodzy, Milind Buddhikot, Tom Rondeau, Keith Nolan, and Doug Sicker.

But mostly, I wish to thank my wife, Rita, for all of the support to this endeavor. In addition to assisting in the research, she made numerous passes through the full length of the dissertation to make discussions more concise, sentences clearer, explanations crisper, graphs more readable, as well as sentences properly punctuated.

Related Publications

Journal Papers

P. Marshall, "Extending the Reach of Cognitive Radio," *Proceedings of the IEEE*, Apr. 2009, vol. 97, No. 4, pp. 612-625.

P. Marshall, "Cognitive Radio as a Mechanism to Manage Front-End Linearity and Dynamic Range," *IEEE Communications Magazine,* Mar. 2009, vol. 47, No. 3, pp. 81-87.

Books

P. Marshall, "Spectrum Awareness," Chapter in *Cognitive Radio Technology,* Elsevier 2007, B. Fette, Ed. Reprinted in *RF and Wireless Technologies*, Newnes, 2008.

P. Marshall, "Spectrum Awareness and Access Considerations," Chapter appearing in *Cognitive Radio Technology,* 2nd Edition, Academic Press, 2009, B. Fette, Ed.

Refereed Conferences

P. Marshall, "Recent Progress in Moving Cognitive Radio and Services to Deployment," in *9th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, Newport Beach, CA, June 2008.

P. Marshall, "From Self-Forming Mobile Networks to Self-Forming Content Networks," in *Association of Computing Machinery Mobile Communications Conference*, San Francisco, CA, Sept. 2008.

P. Marshall, "Dynamic Spectrum Management of Front End Linearity and Dynamic Range," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, 2008.

P. Marshall, "Closed-Form Analysis of Spectrum Characteristics for Cognitive Radio Performance Analysis," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, 2008.

M. Kokar, D. Hillman, S. Li, B. Fette, P. Marshall, M. Cummings, T. Martin, J. Strassner "Towards a Unified Policy Language for Future Communications Networks: A Process," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago IL, 2008.

M. Cummings, S. Li, M. Kokar, B. Fette, B. Lyles, P. Marshall, D. Hillman "Activities of SDR Forum MLM Working Group on a Language for Advanced Communications Systems Applications," in 80th *Software Defined Radio Forum General Meeting*, Boston, MA, 2008.

P. Marshall, "Progress towards Affordable, Dense, and Content Focused Tactical Edge Networks," in *2008 IEEE Military Communications Conference*, San Diego, CA, 2008.

Table of Contents

Table of Figures

Table of Tables

Chapter 1 Quantitative Cognitive Radio Technology Analysis

1.1 Introduction

The objectives of wireless communications have not changed significantly over time. Even before the advent of long distance RF communications, Nicola Tesla foresaw the objective to be:

> "*an inexpensive instrument, not bigger than a watch, will enable its bearer to hear anywhere, on sea or land, music or song, the speech of a political leader, the address of an eminent man of science, or the sermon of an eloquent clergyman, delivered in some other place, however distant. In the same manner any picture, character, drawing, or print can be transferred from one to another place.*" [1]

For most of the history of wireless communications, the primary enabler for enhanced radio communications utility has been the continual improvement of implementing technology. Advances in wireless component technology have enabled fundamental advances, such as transition to narrower signaling, amplifying technology (vacuum tubes), generation of digital and error correcting waveforms, higher frequencies, and lower noise/higher sensitivity receivers. Measurement of these advances, and assessment of the relative advantages and costs of alternative technologies were relatively straightforward and scalar, using measures such as Bit Error Rate (BER), data rate, spectral efficiency, and energy efficiency. These could be compared with the relative costs of incremental technologies in order to optimize cost and performance for a wide variety of applications. This process was possible because the primary measures of effectiveness were directly related to the component technology contained within the radio.

One can envision that advances towards Tesla's goal have occurred, or will occur, through three epochs. In the first epoch, delivery strategies closely paralleled the analog processes (such as FM and AM transmission, where radio link range established the content delivery range). The second epoch has a focus on service architecture, and is characterized by a transition from solutions paralleling analog processes to data delivery strategies that rely on architectural, rather than link performance, incorporating technologies, such as packet switching [2], Internet, and cellular systems. A potential third epoch is the advent of the adaptive and cognitive radios and networks. Cognitive radio technology offers communications whose fundamental advance is a wide range of potential operating modes, and environmentally aware adaptations to select the most appropriate means of communicating data, and cognitive networks adapt their information delivery strategies and topologies to the physical capabilities of the underlying communications links. Where the second epoch was characterized by the use of digital logic to control a stream of data, ***this third epoch utilizes algorithms and experience to determine the strategy for delivering a stream of information***.

Many of the cognitive radio concepts can be considered to be recovering lessons from this first epoch. While modern devices are sophisticated, they lack the human intuition that let High Frequency (HF) operators follow an etiquette of listening for clear channels before transmitting, managing the balance of RF and IF gain to reduce the effects of front-end intermodulation, dynamically selecting frequencies to match the communications footprint, and generally being aware of the communications environment in order to strategize a means to deliver message traffic. Cognitive radio brings the intuitive and environmentally aware principles of the first epoch together with the device and infrastructure technology of the second to produce a new generation of devices that are both sophisticated and clever.

Reed describes the benefits of Software Defined Radio (SDR) as "multi-functionality, global mobility, compactness and power efficiency, ease of manufacture, and ease of upgrade." [3] This dissertation examines the opportunity to use SDR flexibility, plus adaptation mechanisms, to achieve the flexibility, affordability and reliability goals postulated by Tesla.

With the advent of cognitive radio, attention has turned to the behaviors that can be introduced into a wireless device, and the potential advantages of complex, adaptive, and learned responses that could create functional benefits over conventional (static) radios [4], [5]. In parallel, the author proposed [6], [7] that another view of cognitive radio could be that the cognitive features could mitigate the impacts of intrinsic limitations in hardware component performance, and could therefore serve as a method to reduce the cost and energy consumption of wireless devices, through self-adaptation to the environment.

Cognitive radio technology provides the opportunity to forge a unique relationship between a radio and its environment. Cognitive radio adaptation mechanisms will be shown to offer three fundamental advantages:

1. Lower cost, through reduction in hardware performance levels required to address stressing environments,

2. Higher reliability through mitigation of many of the environmental effects that otherwise constrain the reliability of wireless systems,
3. Increased aggregate capability, through denser device and spectrum operation, with mitigation of the undesirable interference effects through adaptation.

However, before any of these propositions can be considered, it is necessary to quantify the answers to two fundamental questions; which are the objectives of this dissertation:

1. What is the degree to which the hardware capability of a cognitive radio can be reduced through reliance on the adaptation of the physical layer?
2. What is the improvement in performance that is achieved by the introduction of one or more cognitive radio adaptations in the decision process of a wireless device?

Much of the cognitive radio research has focused on investigating the potential functions and uses of a cognitive radio. This dissertation starts with an appreciation of the performance impacts of some constraining components within wireless devices, and examines the proposition that cognitive radio adaptations could mitigate these otherwise intractable component shortfalls: thereby, avoiding the necessity of higher cost/higher power consumption designs, without compromising reliability or performance. This is a vision of a cognitive radio that focuses on overcoming component and architecture shortfalls, rather than introducing new user functionality into the device.

This technology has the potential to fundamentally change the way in which wireless architectures deliver services, and changes the nature of wireless communications in a similar way that the advent of microcomputers changed, and largely eliminated the computer time-sharing architectures that preceded them [8]. The range of services, products and economic activity that microcomputer technology initiated was unimaginable when computer services were centralized, and provided largely through infrastructure-based service architectures. The explosive growth of Wi-Fi shows some hint of what technologies based on direct user access to the spectrum could portend, if mechanisms could be devised that provided economic, reliable, and scalable wireless services that were not dependent on infrastructure [9].

In this dissertation, the options for adapting wireless operating modes, and the impact on the reliability and performance of a cognitive radio will be analyzed and characterized in an integrated structure that reflects a range of candidate cognitive functions. Cognitive optimizations adapt performance based on real-time, or near real-time sensed conditions, rather than static, preplanned, or statistical assumptions about future conditions. This is significant for several reasons:

1. Instead of viewing the various cognitive adaptations as independent, they are integrated, and the inherent synergies between them is recognized. For example, spectrum sensing is typically associated with spectrum frequency selection, however, it also offers largely unexplored opportunities to determine pre-selector bands to reduce front-end overload, and to reduce effective receiver noise floor at minimal additional cost to the device.
2. There is no presumption that a cognitive radio is inherently more complex or expensive than its non-cognitive equivalent. On the contrary, this dissertation develops the argument that some of the components implementing a cognitive radio can be less capable in conventional metrics; because the cognitive process can be substituted for traditional approach of achieving link performance and reliability through high levels of device performance.
3. The analysis is generalized and parametric so that the results can be extrapolated to a much wider range of specific environments than are represented by the sample environments available at the time this analysis was performed.

The deployment of Dynamic Spectrum Access (DSA) has been proposed as one of the earliest potential applications of cognitive radio. In this dissertation, the argument is made that DSA can have significant advantages as an enabler of major reductions in the cost of wireless communication equipment through the flexibility and adaptation that DSA provides, independent of the benefits arising through flexible spectrum access. In the framework of this dissertation, the adoption of DSA is argued not from just a spectral management perspective, but as a necessary prerequisite for the introduction of reduced cost wireless devices. DSA is approached from the perspective of maximizing the benefit to the user/owner and operator of a cognitive radio or network through adaptation to the environment, and other users, not just from the perspective of interference avoidance to other users.

As an example, the author, in announcing the Defense Advanced Research Projects Agency (DARPA) Wireless Network after Next Program (WNaN), postulated that adaptive and spectrum aware technology could provide significant relief from high cost and high-energy consumption analog components in receivers[10],[11]. The program announcement hypothesized that highly adaptive network frequency assignments could reduce front-end linearity and Spur Free Dynamic Range (SFDR) requirements by 10 to 20 dB, particularly when supported by new generations of tunable filter technology. Although cognitive radio had been generally viewed as offering advantages in behavior, this announcement asserted that cognitive radio technology offered a path to more affordable and energy conserving devices, through relaxation of component performance requirements. This approach leveraged evolving SDR technology to create a class of devices that could not be designed using legacy approaches to RF design [12], [13].

The cognitive radio construct has the virtue of automating many of the management decisions regarding the device so that:

(1) These operations need not be performed in advance by human designers, planners, or operators, and

(2) They do not need to be based on assumed (rather than actual) environments; can exploit less than worst-case conditions; and can avoid situations that could lead to link failure.

A single radio is not a useful device unless its operation has been appropriately coordinated to establish agreed up technical operating decisions, such as frequency, modulation, and communications protocols. This is nowhere more apparent than in emergency communications, where unplanned, and even inconceivable operational needs are suddenly emergent; resources are inherently strained; communications to rapidly plan and coordinate more communications are difficult or non-existent; and organizations and nations that have never coordinated communications or operations must suddenly do so; generally in the absence of the infrastructure they had been assuming in the past! A good example of these effects was in the recovery operations following Hurricane Katrina [14], or in the Pacific Rim when mobilizing recovery forces after the 2004 Tsunami [15]. A cognitive radio, in addition to enhanced performance, inherently has the ability to make the operational decisions required to establish communications, even in the absence of the preplanning that is integral to current communications operations. A cognitive radio should be capable of operating with only the information required to identify the owner, their information interests, and assess, decide, and then coordinate all other operational communications and networking decisions.

There is a considerable base of prior research literature describing how cognitive radio algorithms could be implemented. For example, work ranging from Mitola [4], [5], to Neel [16] provides analysis of how a network of such devices should behave. While this previous work points out significant advantages to cognitive radio, there is little prior work that measures or predicts the quantified benefits to user and owners through investment in introducing cognitive radio techniques. Even more significantly, there is little consideration or analysis of the ability of cognitive radio technology to provide operation with lower cost or resource devices, by achieving equivalent performance to those operating without the cognitive radio techniques. It is the system performance and affordability benefits that this dissertation will explore and quantify for a range of environments.

The approach to performance analysis will be based on the probability of achieving a given level of performance. While inexpensive consumer communications equipment, such as Wi-Fi cards and Family Radio Service (FRS) radios address reliability as a secondary objective to cost, many communications missions have reliability as the central consideration that drives their cost. A two-megabit per second radio that works every other day is generally less desirable than one that achieves only one megabit per second 100% of the time, even though the mean value of throughput is identical. This is apparent when one contemplates rare, but expensive to mitigate link margin considerations such as rain rate, extreme fading, strong adjacent channel occupancy, etc. Comparison of cognitive and non-cognitive radios will be in terms of the equivalent probability of achieving stated performance thresholds.

The fundamental premise is that a cognitive radio can provide equivalent performance and reliability to a conventional radio, but do so at a lower resource demand (spectrum, energy and acquisition cost) due to relaxation in the component requirements that are necessary to achieve required performance levels. In this framework, cognitive radio technology would not represent additional development and deployment cost. Instead, it would offer the opportunity to reduce wireless equipment costs significantly, potentially changing the fundamental economics of existing and emerging wireless technologies. This is particularly relevant to applications that must provide high reliability. The analysis focuses on confidence in service levels, rather than increase in the mean level of service[1].

In order to fully develop this concept, it will be necessary to fully characterize the different environments in which a cognitive radio would be expected to operate. Conventionally, radios have been built with fixed expectations of the environment they will operate in. Typically, regulators or spectrum managers could ensure these assumptions (such as protected or primary spectrum users) were valid through regulatory or system planning functions. Radios could be designed with a reasonable understanding of the environment they would face. Military radios would face jamming and co-site problems, while consumer cell phones had to receive only one signal from a base station tower, and could filter all other emitters in their vicinity. The effectiveness of a cognitive radio will derive from an ability to operate in the widest range of environments without pre-planned protection mechanisms or strategies. Therefore, this dissertation develops both a specific and generalized understanding of the environments in which a radio must

---

[1] This reflects that, in most cases, bandwidth cannot be "banked", and is not fungible. Excess capacity on one link cannot address reliability on another link, or another time, in most general situations.

contend. Instead of adapting policy to fit the needs of radios and services, it explores how devices can adapt themselves to whatever the prevailing spectrum environment is, and the capabilities of their host (or network partners) hardware. It has been assumed that the first applications of cognitive radio would be for military applications [17] due to their likely high cost: this dissertation will argue that their application to civil needs could and should be as, or more rapid, due to their ability to lower device cost!

In general, this dissertation focuses on non-spread spectrum signaling. The application of spread spectrum to ad-hoc and peer-to-peer communications is constrained by the near/far problem and maintenance of code orthogonality at distributed receivers. Therefore, spread spectrum techniques are not analyzed as suitable candidates for the architectures discussed.

1.2 Structure of this Dissertation

The overall structure of the remaining chapters of the dissertation is shown in Figure 1-1.

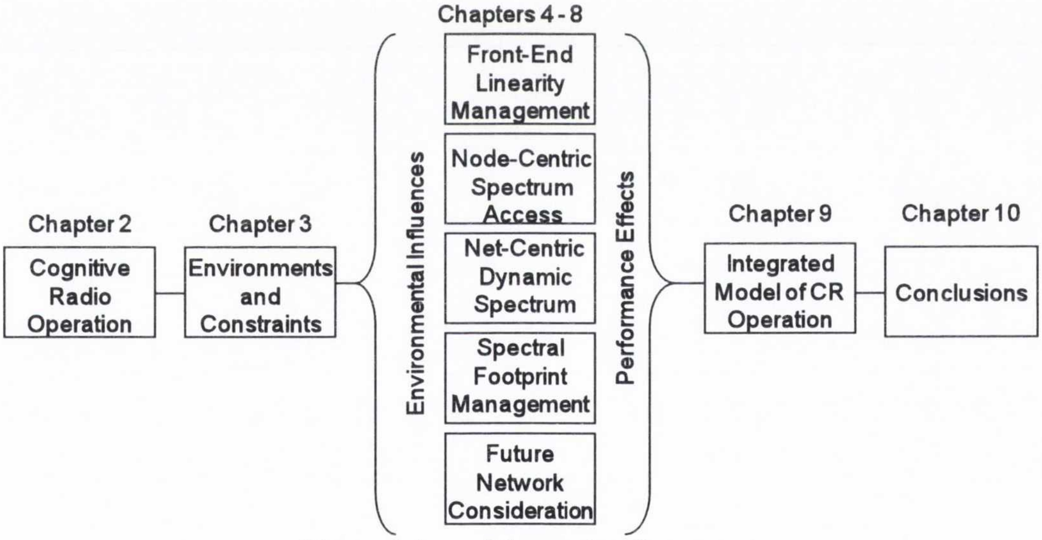

Figure 1-1 Dissertation Structure

This dissertation provides a general discussion of cognitive radio, its operational environment, then individual chapters on specific techniques, and finally, conclusions that integrate the results from the individual chapters.

Chapter 2 presents a general model of a cognitive radio, and describes how its algorithms and decisions interact with both the radio's environment and its operation. There are other algorithms that could be applied, and many implementations may choose to not include, or implement differently, one or more of these algorithms. Since the intent of this dissertation is to determine the availability of acceptable operating points, the actual algorithms used for their selection are not material to the conclusions that are reached.

Chapter 3 examines some of the key environmental characteristics of typical environments, with a focus on those that impact the operation of the physical layer of cognitive radios. It establishes the environmental baseline that will be used to examine the impacts and benefits of cognitive radio technology. Chapter 3 provides measured spectral environment data for environments in which cognitive radios will operate, and develops generalized, closed-form representations of each environment. Prior work [18] has been reported on the organization of spectral samples: in this dissertation, the environmental characterization is a closed-form, probabilistic characterization of the environment, rather than enumerative representation of spectrum data.

Chapters 4 through 7 examine the specific cognitive radio techniques introduced in Chapter 2, and establishes general performance bounds for each technique, using the environments characterized in Chapter 3. These methods include a number of cognitive radio techniques: front-end linearity maintenance; use of DSA; channel selection that minimizes the noise floor (selection among alternative DSA alternatives); and spectral footprint management. Each chapter identifies an environmental consideration, proposes a set of environmental adaptations, and quantifies the benefit to the cognitive radio's capability and reliability. The performance analysis of specific environments is developed in closed-form analysis that enables extension to arbitrary environments.

Chapter 8 introduces network layer processes that can supplement the physical layer performance. These approaches are developed qualitatively. They include network layer or collaborative decisions; such as managing dynamic bandwidth, selection of topology, and content placement. Although all of these techniques operate above the physical layer, they have implications that can impact physical layer decision making, or that are enabled by the physical layer flexibility and awareness afforded by cognitive radio.

Chapter 9 integrates contributions of many of the independent adaptations to develop aggregate performance metrics, and provides the inter-relationship of the decisions across the specific objectives of the individual chapters. Chapter 10 summarizes significant conclusions of the dissertation, and identifies some unmet research challenges that must be explored in order to fully exploit the potential of cognitive radio.

Throughout this dissertation, terms will be used consistent with *IEEE Standard Definitions and Concepts for Dynamic Spectrum Access: Terminology Relating to Emerging Wireless Networks, System Functionality, and Spectrum Management*, IEEE Std 1900.1-2008™ [19] unless otherwise noted.

1.3 General Assessment Methodology and Metrics

There are several general assessment criteria that are common to all of the decisions of a cognitive radio. For each criteria, the benefits will be considered from two perspectives: (1) with performance held constant, how can hardware requirements be reduced; and (2) with hardware capability held constant, how can node or network performance be improved?

The first criterion is the effect of different environmental conditions on the mean or expected value of critical and constraining link performance measures. This measure is appropriate for non-mission critical applications, or in situations where there are alternative paths, methods, or devices that can be applied, or when the role of the device is such that an occasional reduction in capability would be acceptable.

The second criterion is the reliability of meeting critical operational performance measures in a range of environmental characteristics. For many communications systems, the driving requirement is not absolute performance, but assured operation, even in highly stressed environments. This measure assesses how proposed adaptation mechanisms can assure operation through mechanisms other than intrinsic hardware performance. The evaluation matrix is shown Table 1-1. Chapter 2 establishes the metrics by which each of these questions is evaluated.

Table 1-1 Metrics of Cognitive Radio Performance Benefits

	Mean Value of Performance	**Reliability of Operation**
Improvement in Performance	How much more capability would a cognitive radio have, compared to the identical hardware in a non-cognitive radio mode?	How much more reliable would a cognitive radio be, compared to the identical hardware in a non-cognitive radio mode?
Reduction in Required Hardware Capability	How much can the hardware component performance metrics of a cognitive radio be reduced and still achieve equivalent performance levels to a non-cognitive radio?	How much can the hardware component performance metrics of a cognitive radio be reduced and still achieve equivalent reliability levels to a non-cognitive radio?

Chapter 2 Summary of Generalized Cognitive Radio Functionality

This chapter will describe a general concept of operation for a cognitive radio. It will also identify those aspects of a cognitive radio that are relevant to the analysis, and further define that portion of a cognitive radio that is relevant to modifying physical layer performance by adapting the behavior of the lower layers of the device. The description of cognitive radio provided in this chapter will be expanded in later chapters to construct a characterization of the performance implications of the implementation of cognitive radio.

2.1 Prior Work in Cognitive Radio

The term cognitive radio was first used by Mitola [5] describing it as:

> *"the point at which wireless personal digital assistants (PDAs) and the related networks are sufficiently computationally intelligent about radio resources and related computer-to-computer communications to: (a) detect user communications needs as a function of use context, and (b) to provide radio resources and wireless services most appropriate to those needs."*

Mitola references both ontological and learning mechanisms for the computational intelligence. Potential benefits of this intelligence were in the areas of "Radio Resource Management, Network Management, Service Delivery, and Type Certification." Of the "radio resources" of interest, spectrum pooling[1] is cited as one of the most significant.

Since the Mitola work, a number of different approaches to, and definitions of, cognitive radio have appeared. For this dissertation, the working definition of a cognitive radio is a radio that is aware of its environment, and constantly utilizes this awareness to adjust its operating characteristics and behaviors to provide effective performance in a wide range of environments. In contrast, a conventional or non-cognitive radio typically has these same decisions made during the design process, or at least in advance of its operation, based on assumptions about the likely future environment (such as dynamic range, spectrum occupancy, path characteristics, etc.). Haykin [20] provides a concise list of cognitive radio characteristics:

1. To perceive the radio environment (i.e., outside world) by empowering each user's receiver to sense the surrounding environment continuously.

---

[1] Spectrum pooling is defined as a number of applications, users, or networks sharing a fixed set of frequencies and using them on an "as-needed" basis, based on the statistical demand.

2. To learn from the environment and adapt to it in response to deviations in the environment.
3. To facilitate communication among multiple users through co-operation in a self-organized manner.
4. To control the communication resources among the multiple users through competition.
5. To create the experience of intention and self-awareness.

The first three of these items are the focus of this dissertation; the methods of control and the type of experience will not be addressed.

The recent upsurge in interest in cognitive radio technology attests to the rapid growth in understanding and appreciation of its potential to revolutionize the interaction between a radio, its environment and its user. While it has been known that the environment has strong (and usually deleterious) effects on the operation of a radio; only recently has the capability emerged to exploit awareness of the radio environment. Traditionally, radio designers have considered how a fixed mode of operation is degraded in a variety of less than ideal situations. Cognitive radio now presents the opportunity to investigate adaptive modes to maintain and optimize performance across a range of actual environments.

It is instructive to compare the cognitive radio algorithm design process to that of a conventional radio. In conventional radio and link design, the required link reliability drives the degree to which unlikely environmental effects are considered and accommodated in the design. The higher the required reliability, the more design requirements must be driven by worst-case assumptions regarding factors such as rain and ice attenuation, in-band noise, adjacent channel energy, and propagation. While these designs can yield high reliability communications; when actual conditions are less stressing, the device often cannot exploit the opportunities resulting from the less than worst-case conditions. A measure of the effectiveness of a cognitive radio is how closely it can match its performance to actual conditions and usage. This supports the concept that reliability should be one of the fundamental motivations for cognitive radio.

Previously, design studies stated an assumed environment for which performance was optimized: a cognitive radio design instead can argue for a range of environments, each with individually and continually optimized performance. This enables radio operation to transition from being determined by the radio designer, based on a projected environment, to being dynamically determined by the radio and/or network, based on a continually maintained perception of its environment. This can be approached as an opportunity to achieve incremental gains in radio performance, or as an opportunity to fundamentally change the relationship of radios, their environment, and the networks and applications that run over them.

The transition of cognitive radio from concept to reality has made considerable progress in the last several years, and has been largely pursued on two paths. The first is the drive to mature, demonstrate, and to ultimately deploy, DSA systems [21]-[27]. A number of commercial and governmental organizations have completed prototype implementations and submitted them for varying degrees of technical or operational scrutiny. This effort broadly clusters into two application approaches: with much of commercial and industry focused on unlicensed access to television White Space, such as contemplated by the IEEE 802.22 Wireless Regional Area Networks [28] -[30], and research into more broad-based spectrum sharing, typically focused on existing licensed applications, such as military and public safety [31]-[37] as the initial applications.

Progress specific to 802.22 applications has been driven by the opportunities provided by technical achievements and potential regulatory acceptance of television "White Space" utilization [38], [39]. Published work demonstrates that significant spectral and geographic regions exist that could support "white space" operation with feasible detection thresholds [40]. Proposed approaches vary from use of centralized regulatory databases, to fully distributed and fused sensing. Of particular interest is the reliance on multiple resolutions of spectrum sensing to control the operation of the device, with both fast (1 ms) and slow sensing cycles (25 ms) driving the adaptation of the access point and customer equipment. Implementations of receivers supporting this architecture have been reported [41], supporting the technical feasibility of this approach. One approach relies upon analog spectral analysis and temporal correlation, with only the correlation difference requiring digitization at a very low rate, a departure from the assumption of all-digital sensing [42]. Increased processing in the analog domain, and processing of temporal signatures promises much lower energy consumption and potentially lower cost than high sensitivity and high-speed digital approaches. Also, it has been shown that cyclo-stationary sensing for incumbent sensing in poor and fading conditions is feasible with limited computational resources [43]. Scheduling of cognitive radio sensing within and across communities has been explored and shown to be a feasible approach to the implementation of multi-domain cognitive radio operation [44].

While much research and discussion has focused on the "White Space" arguments, important, and relatively neglected effects of adjacent channel energy are beginning to be discussed within the framework of Cognitive Radio. A simple "occupied/not-occupied" categorization has been the basis of much of the discussion of spectrum availability. Consideration of adjacent band effects [45] shows that secondary use of television frequencies in adjacent bands has minimal effect on television reception, and that reasonable performance and interference avoidance is achievable at relatively close distance to an adjacent channel television transmitter.

Perhaps in recognition of both the likely deployment of sensing based systems, and the history of malicious applications of technology, methods of detection and mitigation of falsified incumbent signals, and assurance of trust in distributed cognitive radio sensing data has now been proposed. Trust assurance must be an essential element of the research agenda, if the sensing scope of Cognitive Radios is to transcend individual trust domains [46], [47]. Without these features, it is hard to imagine how any user or operator would be willing to trust the operation of their network to a technology so vulnerable to denial of service (DOS). The second major research topic is the increased understanding and initial prototyping of learning or genetic based cognitive radio algorithms. Work reported by a number of researchers has demonstrated performance levels that should be sufficient to make the argument for incorporation into operational radios in the near future. [48]-[58].

Early work on the representation of knowledge within a cognitive radio has also been extended. Yarkin and Arslan [59] describe a Radio Knowledge Representation Language (RKRL) that provides a mechanism to organize multiple sources of environmental and location awareness into an integrated representation. Of particular importance is that this approach provides a method to organize the channel awareness into a structure upon which a cognitive radio can make operating decisions. In a similar vein, other researchers have extended knowledge representation and encapsulation to describe the state of knowledge of the radio components of the transceiver [60]. Ultimately, such understanding must extend beyond the wireless device itself, and extend to the network and its environment, and users [61], [62].

Not only has research in DSA advanced the technology over the last several years, the regulatory and spectrum community has become increasingly involved in the process, evidenced by the success of the policy program at the three IEEE Conferences on New Frontiers in Dynamic Spectrum Access Networks (DYSPAN). Corresponding interest in industry and regulatory standards is reflected in transition of IEEE P1900 to a Standards

Coordinating Committee for Dynamic Spectrum Access Networks (SCC41) [63], [64]. The breadth and international scope of membership offer the opportunity to advance DSA as an integrated package of technologies, rather than as individual and standalone efforts addressed on a nation-by-nation basis.

Several complete purpose-built Cognitive Radio experimental platforms have been demonstrated, or are under development. The Kansas University Agile Radio, sponsored by the National Science Foundation (NSF), appeared as one of the platform alternatives on which Cognitive Radio waveforms and algorithms can be constructed and tested without being constrained to the characteristics of the current commercial waveforms and the embedded lower layer behaviors [65]. Nolan et al. report joint experiments with two different cognitive radio systems to demonstrate detection, coexistence, and interoperability [66]. A complete end-to-end framework of a cognitive radio is provided by Raychaudhuri et al. [67].

DARPA has initiated a "second generation" purpose-built cognitive radio program [10]. This program, WNaN, uses DSA as its fundamental operating principle, and trades high performance individual transceivers for replicated, but lower performance transceivers. The operating principle would be adapting around situations that would otherwise cause front-end overload, as will be discussed later in this dissertation. One of the stated objectives of this program was to demonstrate that cognitive radio would produce at least the performance of a conventional radio, but accomplish this with reduced performance components through use of adaptation to mitigate the performance stressing environments that would otherwise drive the design to performance that would increase energy consumption and likely cost. The quantitative basis for front-end performance adaptation is provided by Marshall [7]. The WNaN radio has four independent transceivers, covers 900 MHz to 6 GHz, has high quality front-end filters to enable the DSA functions to identify and utilize low energy pre-selector bands, and uses a commercial Radio Frequency Integrated Circuit (RFIC) for the transceiver functionality. This platform is intended to provide not only DSA functionality, but cognitive topology management, and content management, as well.

The discussion in this dissertation assumes that the device or network is authorized by some process to enter the spectrum and use it, so long as its operation is consistent with a set of policies. Considerable research has gone into the subject of how market and allocation mechanisms can arbitrate allocation of spectrum through centralized and distributed algorithms. Yuan et al. [68] describe a typical structure capable of both central and distributed spectrum allocation based on pair-wise interference. Zhao et al. [69] describe a purely distributed process for arriving at coordinated spectrum allocations. Both results demonstrate that the introduction of centralized or distributed coordination does not fundamentally change the process, or result in a likelihood of far from optimal solutions.

Most DSA research has focused on ensuring non-interference to other spectrum users. In this dissertation, the technology to ensure non-interfering operation is assumed to be sufficient. This dissertation investigates the question of how a radio can determine the relative advantages of each spectrum segment for which it is provided a policy compliant opportunity. Not all spectrum choices are equal, and the desirability of a given piece of spectrum is a function of more than just the channel's characteristics. From the node or link perspective, it includes the effects of adjacent band usage, as perceived by the network members, the likelihood of forced relocation, the predicted noise floor, and the denial of spectrum access to other network nodes. From the network perspective, it includes the effect of decisions made by one node on the operation of all other nodes sharing that spectrum. In this discussion, spectrum usage is considered and optimized from a "bandwidth over geographic area" perspective, rather than optimized to reduce the occupied bandwidth of the signal. In this dissertation, DSA is an enabler to optimize link and network operation, not an objective, in and of itself.

This process is independent of the mechanism of spectrum sharing[1]. With knowledge of the impacts of spectrum decisions, the device or network can adjust its own independently determined choices, make requests to spectrum brokers, or even make bids into an auction process to best reflect the utility of each spectrum alternative.

2.2 Cognitive Radio Functional and Performance Model

Since the term "cognitive radio" will be used often, it is important to examine its possible definitions. There appear to be two fundamental definitions that should be considered. When Mitola [5] first referred to the term, he was referring to a device that had the characteristics of an intellectual process; implicitly a learning system, such as envisioned by decades of Artificial Intelligence (AI) research. A (perhaps oversimplified) litmus test of such systems might be that they derive their behaviors from experience-based learning.

---

[1] Such as decentralized, micro-charged, brokered, auctioned, etc.

Consistent with that definition of cognitive radio, researchers have reported progress in applying learning engines to the selection of channels and waveforms [70].

In usage, the term cognitive radio has taken on a significantly broader connotation, and less technologically challenging definition. The DARPA Next Generation Communications (XG) program reported success with DSA using fixed algorithms and policies [71]. The United States (US) Federal Communications Commission (FCC) undertook a cognitive radio access proceeding [72] focusing specifically on dynamic spectrum behavior, and was silent on how such behavior would be implemented. Several industry organizations have announced progress on cognitive radios, although, there is little published information on that the designs, algorithms, behaviors, and most importantly, performance.

Although the strict definition of "*Cognitive*" might not seem to include deterministic or declarative technologies; if we wish to assert that cognitive radio implements "cognition", a fundamental distinction between a cognitive radio and a non-cognitive radio would be the degree to which the device makes its decisions based on awareness of the environment. Such a definition has advantages since it is focused on "*What*" the device does, not "*How*" it does it.

The necessity to be open to both definitions is apparent in the area of frequency selection: two different and competing concerns are being addressed through cognitive radio; simultaneously achieving both non-interfering and optimized operation. An insistence that a cognitive radio must derive its algorithms through learned behavior provides no mechanism to address the inherent constraint that the effect of some decisions is only available external to the learning system itself. In this case, the necessary feedback to reinforce behavior is not available. The impressive results reported for physical layer learning are achievable because the two link partners constitute a closed system that can use itself to provide reinforcement to the learning. In contrast, a system comprising many heterogeneous users of spectrum, with incompatible waveforms, protocols and even operational concepts[1], then it is much harder to apply feedback-based learning as the sole mechanism to discover and correct "*aberrant*" behavior.

There are other practical issues that should be considered. Given the poor state of Information Assurance (IA) technology, few system owners or implementers would be willing to have their system depend on information received from, or requiring exchange

---

[1] For example, a broadcast, receive only node has no ability to report interference events.

with, an external party. The growth of Internet malware demonstrates that it is unrealistic to require participants to cooperate and exchange information that will impact behavior and performance. Engineers might argue that the exchanges were beneficial, or at least harmless, but the seemingly endless discovery of exploits and unintentional behaviors in commercial software products would undoubtedly hinder any solution that required nodes to trust actions or information provided by nodes external to their own trust domain.

Regulators tightly control spectrum, due to the extreme social and economic consequences of interference to key infrastructure and services. It is unlikely that they would accept the principle that although a nodes behavior will initially be very poor, it will eventually learn to "*play nice*" in the spectrum (after causing any number of interference events while learning!). Similarly, they might not be pleased to learn a device tested and accepted as behaving well and not interfering would have the capability to later learn "bad habits" that would have the very consequences the technology was intended to avoid! Their awareness of the potential for this behavior is evident in expressed concerns, such as by the UK's Office of Communications (OFCOM) [73].

There is a strong argument that when improper behavior has negative impacts on other, external spectrum users, than the use of declarative techniques is advantageous. These techniques offer an approach that can be analytically proven, experimented with to any desired level of confidence, and then "frozen". Admittedly, this is not as intellectually pleasing as a radio that "programs itself", but it is testable, verifiable, predictable, and perhaps even provable.

This provides a clear partitioning of the cognitive radio reasoning infrastructure. On one side, there are functional behaviors such as sharing and etiquette that focus on the device's impact on other users in the environment. Presumably, this behavior is externally mandated, as the radio does not benefit from the constraints, and might have its performance reduced. On the other side, there are adaptation mechanisms that can optimize performance, perhaps within a region constrained by external policies. Not surprisingly, these two research areas have unique approaches to exploiting environmental awareness. The externally imposed constraints may limit the range of potential optimizing strategies; therefore, they are performance neutral, at best. In fact, it is likely that desirable operating modes may exist and be exploitable by the node, but are precluded by external constraints, due to their impact on other users of the spectrum. This structure enables the optimization engine to be focused solely on the best choices for the node, and can assume that other users are protected by policy constraints.

The DARPA XG program has reported that one of its objectives was developing a declarative language and processing tool set for external policy control of spectrum sharing. The program has developed two radio-borne reasoning engines that provide inferential processing of predicate calculus policies based on a spectrum management ontology and rule set [74]-[76][1].

It is appropriate to partition the intelligence model of a cognitive radio into endogenous and exogenous components. While examining structures for policy enforcement and optimization of DSA, the author partitioned DSA cognitive radio decision making into two elements: a System Strategy Reasoner (SSR) to optimize the decisions and performance of the device, and a Policy Conformance Reasoner (PCR) to enforce externally provided policies, as shown in Figure 2-1 [77].

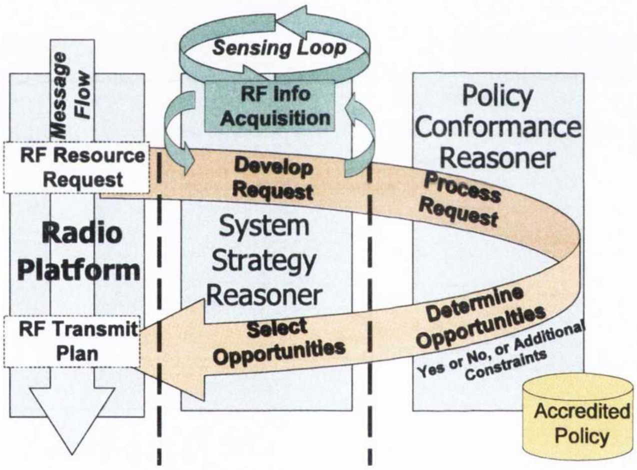

Figure 2-1 PCR and SSR Model of a DSA Cognitive Radio

The PCR provides the external (exogenous) policy component(s) that address the device's impact on the external environment: primarily avoidance of spectrum interference to other users. The SSR provides the internal optimizing behaviors (endogenous) to maximize performance through selection of operating mode and other parameters. This partition reflects the fact that the external effects of the device are difficult, and often impossible to perceive from the reference frame of the individual radio. It is therefore difficult to argue for an exclusively learning model of a cognitive radio, or one based solely on performance optimizing algorithms. Such methods will be oblivious to their impact on other users, and thus, they will not "learn" how to avoid interference to other users, even while successfully learning to avoid conditions that would result in interference to their own operation.

---

[1] The progress to date in the application of predicate calculus reasoners is discussed in more depth in Chapter 9.

An example of this partitioning is apparent in the three principles for non-interfering DSA operation which the author proposed as the standard for evaluating the performance of the XG DSA technology [78]. These were:

1. ***Do No Harm (to other users of the spectrum),***
2. ***Add Value (to the user, operator, or owner who invested in the technology), and***
3. ***Perform (robustly and reliably in a range of environments and user mission needs).***

These principles reflect the role of the cognitive radio intelligence mechanisms in protecting other users from the effect of the device's behavior, and in providing a measurable return on the recurring and non-recurring investment in development and manufacture. These principles address unique and mostly non-overlapping objectives and responsibilities. The first represents the exogenous concerns. The second and third items represent the endogenous objectives. They would appear to be difficult to accommodate in any single abstraction (learning vs. declarative knowledge).

Looking beyond cognitive spectrum management, there is a need to extend the concepts of cognitive radio from a spectrum and link focus, to an integrated abstraction that is inclusive of all aspects of wireless device decision-making [79]. This includes multiple decision domains, such as network routing, topology, user interaction, security, and content management in a structure capable of expressing both constraining and optimizing policies, such as those shown in Figure 2-2. The need for such structures has been recognized in the literature [80]-[82].

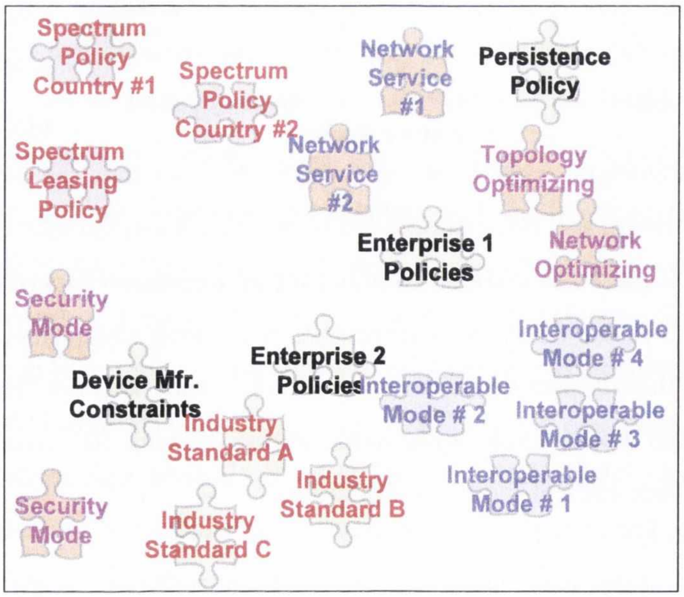

Figure 2-2 Example Scope of Cognitive Reasoning Policies

Two issues must be addressed to make this vision a reality: First; an abstraction of how different cognitive adaptations interact; and second, a computer science implementation of these techniques that is consistent with the resource limitations of typical mobile devices. Without solutions that meet these objectives, the cognitive radio community will be building hundreds of standalone experiments, but few usable products that can benefit from more than one breakthrough due to lack of an integrated mechanism that can unify multiple technologies into a single cognitive radio. As a user of AI technology, the cognitive radio community may be one of the first communities to be able to articulate the need for a unified theory of cognitive automata that can integrate learning and declarative knowledge into a single logical framework, rather than distributed across two disparate mechanisms.

An example of where this technology could be applied is spectrum decision making. Traditionally, the societal interest in avoiding interference to other users of the spectrum is reflected in the direction and policy provided by national regulatory administrations. In the cognitive radio model, these same protective interests are represented by the constraining exogenous policy component. The internal optimization of performance is provided by the endogenous component. The exogenous component(s) ensures that the decisions of the endogenous component(s) do not have unacceptably detrimental effects on other users of the environment with which the cognitive device interacts, such as other spectrum or network users, or the device's own user. This partitioning allows the endogenous component(s) to operate as an unconstrained "greedy" algorithm, since the "social" values of the community (avoiding impact on other spectrum users, network users, or host device) are protected by the exogenous component(s).

This discussion points to the necessity that the advancement and deployment of cognitive radio not be viewed as solely a technical challenge. New and dynamic spectrum regulatory regimes are central to almost all descriptions of cognitive radio. Some degree of spectrum selection and control, with environmentally aware selection among spectrum and waveform options is clearly one of the defining aspects of cognitive radio technology. To achieve this, it is desirable (inevitable) that the regulatory community be highly involved in the development of the technology, and that the technology include features that address regulatory considerations to at least the rigor provided through current manual technical, operational, and regulatory methods [83].

The operation of these processes is depicted in Figure 2-3. The SSR attempts to maximize the effectiveness of the device, given the environment that is sensed, and the needs of the device's current mission and loading. The PCR(s) validates that each action the radio proposes to take is acceptable based on the policy enforced by the device. An external party, such as a spectrum regulator, network owner, or hardware developer can, through a common policy language, provide the actual logic of these reasoners. In a heterogeneous network, a desirable attribute of the policy process is that it specifically enables operation using policies derived from multiple sources.

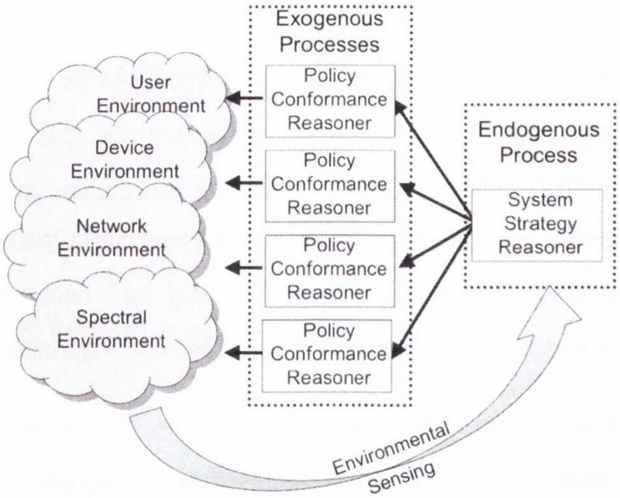

Figure 2-3 Relationship of Endogenous and Exogenous Reasoning in Cognitive Radio

Equally important, these sources of policy should be able to develop policies that are independent of the policy context within the device (i.e., other policies a user may have loaded)[1]. A spectrum owner might provide policies for sharing a band, but these policies would have to coexist harmoniously with policies independently provided by national regulators, and perhaps device limitations provided by the manufacturer. Similarly, a device manufacturer might have products that could be running a variety of cognitive radio software, and would provide policies to ensure that the operating limits of its specific device were not exceeded.

Figure 2-2 illustrates the range of policy domains, and thus necessity of fitting policies together dynamically into the architecture depicted in Figure 2-3. This figure depicts the control cycle of a cognitive radio; sensing the environment(s); creating responsive strategies; ensuring their compliance with exogenous constraints; and (implicitly) implementing them within the device. Some authors have treated the spectrum environment as unique from the user behavior and device status; however, a more general framework treats these as symmetric domains that coexist, and are integrated into the behavior of the cognitive device through mechanisms within the SSR. Intuitively, it is reasonable that as more of these domains are integrated into the option consideration and decision structure provided by the SSR, the more globally optimal the decisions will become. Although the figure shows PCRs for each of the environmental domains, this one-to-one relationship is not necessary if the policy language is general enough. The list of environmental domains is not intended to be exhaustive; but to show the potential scope of a symmetric solution to cognitive radio conformance policies.

---

[1] This is in contrast to the IETF Policy Language, where the position of statements controls their context, due to its procedural nature.

Some illustrative spectrum domain considerations are provided in Table 2-1. They are partitioned between the exogenous and endogenous reasoning processes. In general, it is apparent that the exogenous concerns are those typically associated with communications device regulation; while those of the endogenous component are typically associated with conventional communications link design. This reflects that:

> ***The cognitive radio strategy development is really an analog of the manual design process; different in that it is interactive with its environment; based on measured rather than predicted data, and continually adjusted.***

The design and environmental considerations are the same as for a non-cognitive device.

Table 2-1 Considerations for Endogenous and Exogenous Reasoners

Exogenous Reasoning Considerations	**Endogenous Reasoning Considerations**
Co-Channel Interference	Operation in a Multi-Path Environment
Adjacent Channel Interference	Coding and Error Control Strategies
Interference to Safety of Life Systems	Overall Energy Consumption and Transmit Power Management
Unintended/Spurious/Out of Band Emissions	Waveform Performance in a given Environment
Receiver Performance Assumptions	Frequency Selective and Multi-Path Effects

DSA may be the first domain in which one might argue for this adaptation, but clearly is not the only one; and perhaps, in retrospect, it may not even be the most important. As an example, Issariyakul et al. report that there are significant benefits to control of cognitive radio "greed" regarding radio resources from the perspective of efficient Transport layer operation [84].

Most consideration of the exogenous PCR has been in the spectrum domain and has been focused on the issues associated with spectrum interference and access. The concept is equally applicable to the other environments of the cognitive radio. For example, a network owner might require the device to validate its admission and packet traffic to ensure that the actions of the device were consistent with the network's operating rules, or a device manufacturer might include a validation engine at the hardware interface to ensure that the operation of cognitive radio algorithms did not violate any of the operating limitations of the device (such as thermal limits, transmit duty cycle, amplifier bias, etc.). A conventional network firewall is an obvious specialized implementation of this structure's PCR construct. Similarly, the V-Chip (television programming content control) provides rudimentary video policy enforcement, and Internet content filters are also examples of PCRs that operate at the user interface to ensure that information over this interface complies with an externally provided policy. Thus, the PCR is not a new concept; it is a generalization of functionality that experience has shown is needed to buffer the actions of a technology and the environment it impacts.

Another consideration is the location of these components. Although there is a tendency to think of cognitive radios as peers and self-contained; there is no necessity that implementation of these components be symmetric or local. Spectrum selection and conformance checking can certainly be performed at a central node that is advantaged; network policies can be provided at boundary devices, etc.

Table 2-2 provides some illustrative environments that impact decision making of the endogenous component, the characteristic elements of each environment, and examples of the decision-making objectives. These are decisions that the radio can make without policy constraints due to impacts on other members of the environment.

Table 2-2 Illustrative Endogenous Cognitive Radio Operational Environments

Environment	Characteristics	Significance
Spectral and Channel	Frequency Utilization	Selection of frequency for interference free operation of the selected link frequency
	Energy Distribution	Selection of bands of operation to avoid adjacent channel effects
	Multipath and Frequency Selective	Selection of waveforms to mitigate multipath effects
Network	Traffic Volume and Character	Transport layer tuning and adaptation
	Traffic QOS Needs	Prioritization of packets; latency management
Internal	Energy Limits	Selecting Power Conserving Modes and Low Energy Waveforms
	Thermal State	Imposing Duty Cycle Limitations and Selecting Efficient Amplifier Modes
User	Patterns of Use	Organizing Menus; Sequencing options
	Identity	Pre-adapting to User needs and patterns

The partitioning of the reasoning into these two categories provides a simplifying abstraction. At best, the exogenous policy conformance process cannot improve performance; its effect is only to preclude certain actions. Therefore, it can restrict, or reduce performance, but cannot improve it. The endogenous component is the only one of the two reasoning processes that can contribute to performance, and therefore, it is the primary focus of this dissertation. Neel [16] provides an argument that the behavior of spectrum interaction can become stable, even given locally optimizing behaviors, but this analysis must be extended to include the interaction of these processes with the other layers, and the effects between them.

In summary, the use of exogenous and endogenous reasoning provides a partition between the aspects of cognitive radio operation that address the need to operate within a bound of acceptable external impacts, and those that address performance optimizing by the device(s), based on the objectives of the user. Since the endogenous elements are the ones that optimize the performance of the device, the rest of this dissertation will focus on these processes, with the assumption that the constraints introduced by the exogenous considerations cannot introduce additional optimal solutions. However, throughout this analysis it will be recognized that the exogenous processes may preclude use of otherwise optimal solutions.

2.3 A General Approach to Cognitive Radio Operation

There are a large number of ways in which the processing of a cognitive radio could be organized, and most of this dissertation is agnostic to the actual organization of the processes. However, there are dependencies among decisions and it is useful to have a notional concept of the overall decision framework of the radio, in order to position the individual optimizing process within an overall framework.

Chapters 4 through 7 provide the tactical and quantitative aspects of these processes. Chapter 8 provides qualitative discussion of network level extensions of these techniques, and Chapters 9 and 10 unify the discussion of individual tactical decisions, and develops the synergies among them.

In the structure to be described, sensing is performed at the physical layer, with no cooperation between the service that is assured with interference-free operation (typically, a primary user), and the opportunistic service. However, the same principles are applicable to hybrid architectures, such as proposed by Zhao et al. [85].

A notional, top-level process for physical layer optimizing by a cognitive radio is depicted in Figure 2-4 below. The chart depicts two key loops in the physical layer processing; one selecting low total energy bands with non-interfering frequency alternatives, and one evaluating the available frequencies to select the one that appears to offer the most effective communications alternative.

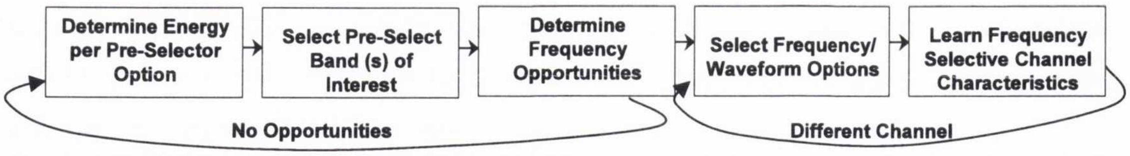

Figure 2-4 Top Level Cognitive Radio Physical Layer Optimization Process

A similar process exists for the selection of higher-level functions, such as within the MAC layer, and selection of network layer routing. These decisions may interact with the physical layer process described in this dissertation. The general principles described in this dissertation are applicable to these additional environmental domains; extending the scope of the evaluation and ranking functions to be described later in this Chapter.

There are any number of ways in which these operations may be implemented and arranged. The specific decisions outlined below may be performed in a series, as described, or in globally optimized algorithms that simultaneously consider the convolution of all possible decisions. Table 2-3 describes each of these top-level processes, and the chapter(s) where they are discussed. For simplicity, the text will refer to these as decisions. The radio could implement the process by maintaining a list of all possible (discrete) options and refining this list successively, and discarding obviously non-viable approaches.

Table 2-3 Notional Description of Key Cognitive Radio Optimizing Processes

Function	**Description**	**Chap**
Determine Energy per Pre-Selector Option	Survey Spectrum to determine the total energy that is present in each pre-selection option, and determine the response of the radio to that energy.	4
Select Pre-Select Band(s) of Interest	Select a specific pre-selector option for each potential operating frequency. This may be as simple as selecting a filter, varying the Q, switching of filters in/out to improve noise temperature, etc. This step ensures that the radio is capable of operation with the selected band, within an acceptable margin of intermodulation energy.	4
Determine Frequency Opportunities	Determine one or more frequencies within a selected band or set of bands in which the radio can operate consistently with the interference policies that are enforced.	5, 6
Select Frequency/ Waveform/ Modulation Options	Select a specific frequency of operation, modulation, and power level or spectral distribution. These choices are interrelated, since the size of the spectral openings may constrain the choice, waveform modes that are optimal for reception may be less optimal for signal generation, different susceptibility to adjacent channel signals or other considerations.	7
Optimize Network Level Decisions	Consider the effects of error rates, topology, bandwidth usage and content access patterns in link and topology decisions.	8

At first appearance, the process depicted above appears to be one that selects a discrete frequency and waveform for operation. A more flexible way to envision this is that the radio is constantly maintaining a list of emission options, and constantly grading them. When the current one is no longer valid, or sufficiently effective in relation to the optimal operating point(s); it is replaced from the list and the new one implemented. When the current one is less desirable than another candidate by a significant enough margin (to be worthy of the investment to change the operating mode), the mode is changed. If the radio has the capability to use discontiguous waveforms, or multiple transceivers, it is not even necessary to view the selection as a down-select to a unitary option. On this case, multiple emission options may be used simultaneously on single or multiple data streams.

Table 2-4 illustrates the performance metrics that will be considered in this dissertation.

Table 2-4 Performance Metrics and Associated Functional Areas

Performance Metrics	Driven By Functional Characteristic
Operational Availability	Reduction in the probability of front-end third–order intermodulation overload through selection of pre-selector bands containing low energy levels. Reduction in probability of front-end third–order intermodulation noise level through selection of pre-selector bands containing low energy levels. Reduction of the probability that a single source will create interference to the communications channel. Maximizing the probability that DSA algorithm will be able to locate an open (unused) channel that does not interfere with other users.
Link Capacity through Reduced Noise Floor	Increased mean throughput due to reduction in the probability of front-end third–order intermodulation noise level through selection of pre-selector bands containing low energy levels. Selection of channels with the lowest in-band noise level.
Increased Operating Period/Reduced Energy Storage	Reduction in probability of front-end third–order intermodulation Noise Level through selection of pre-selector bands containing low energy levels. Selection of channels with the lowest noise level at both ends of the link, or the best available compromise for all link members.
Node Density & Capacity	Increase in the density of nodes using the spectrum through adaptation of frequency selection to minimize the separation of cognitive and non-cognitive nodes. Increasing the density of cognitive radio nodes by optimizing the amount of interference caused between cognitive nodes. Maximizing the information flow in a geographic and spectral extent, through joint optimization of spectrum reuse and information rate.

Zhao et al. points out the difficulty in establishing metrics for cognitive radios and networks [86]. In this dissertation, the intent is not to assess cognitive radio metrics as standalone measures. Instead, they are examined in the context of the performance and the reliability they achieve, and the opportunity they present to reduce hardware performance without unacceptable performance and reliability impacts.

Additionally, many of these performance metrics are impacted by multiple functional areas (as described in Chapters 4 through 8). In this analysis, it is assumed that a cognitive radio is aware of its Signal to Noise Ratio (SNR), and operates at a constant BER. The opportunity provided from reduction in the effective noise floor is modeled as increments in required power, range, or bandwidth.

2.4 Band, Frequency, and Emission Characterization and Selection

The next section outlines a hierarchical process to model the decision making that starts with a range of possible bands and frequencies, and eventually decides on one or more specific frequency, bandwidth and emission mode(s). This framework can be considered to be a subset of a more general model of cross layer optimization within cognitive radio.

Each option available to a cognitive radio can be modeled as a decision option chain, initially routed in available spectrum options, with (potentially multiple) options for utilizing this spectrum through waveform choices, upper layer operation, and other operating options that the cognitive radio system and network provides. Figure 2-5 illustrates a potential hierarchy of spectrum options, each of which, in turn, could be utilized by one or more waveform alternatives, each of which would implicitly represent the choice of a specific network topology. Topology choices are not perfect trees, but may be implemented through a number of permutations of choices. In addition, these choices are not exclusive; a given node could select multiple blocks of spectrum through either discontiguous waveforms, or independent transceiver operation. The topology aspect is important; the purpose of network radios is not to form links, but to move information across networks of nodes. While the early research and experimentation in self-forming and Ad-Hoc networks has focused on node discovery and connectivity, it is the character and performance of the network that influences user experience. This is an important challenge, but as these networks become denser and more expansive, the challenge shifts from forming connectivity, to managing it.

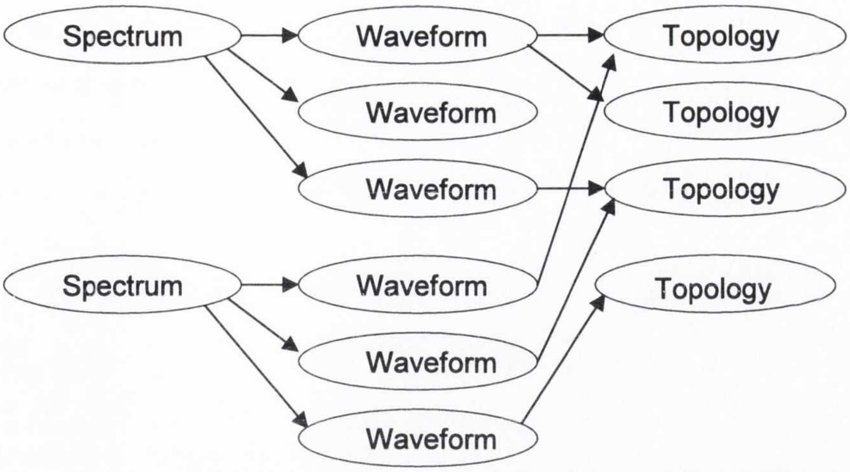

Figure 2-5 Hierarchy Tree of Cognitive Radio Selection Alternatives

Spectrum is an important part of this decision process, since the selection of frequency has an important impact on operating and interference range. Additionally, as a practical matter, link throughput is also typically a function of frequency, as higher frequencies generally have more bandwidth available; technically, due to the increased availability of signaling bandwidth within the fractional bandwidth of the filters and antenna, and, as a regulatory impact, due to the general crowding of lower frequency spectrum. The topology layer represents the network's strategy to move information, given the constraints of waveforms and available spectrum. This aspect of DSA decision-making has received relatively little research attention, although it is a primary objective of the WNaN program [87].

Information on each of these considerations can be represented by a number of different structures, such as formal parameters, XML, or ontology expressions. The use of an ontology as a representation of the exchange between cognitive radio components is appropriate due to their ability to associate the semantics of the description of information with inferential reasoning, and to enable continual evolution of semantics as technology advances, without the necessity to change the underlying ontology. Some structures and representative ontologies have been proposed in the literature by Marshall [88], [89], Denker, Perich, and Elenius [74], [75], [90], and Kokar [91]. The discussion in this dissertation is not sensitive to the mechanism for this exchange.

Figure 2-6 shows that the ***Spectrum*** object can be as simple as a single center frequency, a bandwidth, and power, or it can be as complex as a piecewise described, discontiguos set of variable power levels. In effect, it describes the constraints on the spectral density function of the resulting waveform. At this point in the process, this is strictly a frequency domain expression; its abstraction both defers and hides the details of the waveform from the interference analysis implicit in policy conformance. A precedent of this structure is the US FCC's adoption of the Ultra-Wide Band (UWB) emission mask as the basis for licensing UWB devices in the US [92]. If additional considerations became critical to interference analysis, these would be candidates for extension of the ontology. The important feature of this abstraction is to isolate the constraints that derive from the exogenous concerns about interference to external parties from the endogenous optimization activities that are internal to the cognitive radios operation. The *Spectral Specification* expression is therefore a "hunting license", with which the endogenous optimization functions of the radio can proceed without further consideration of the environmental effects.

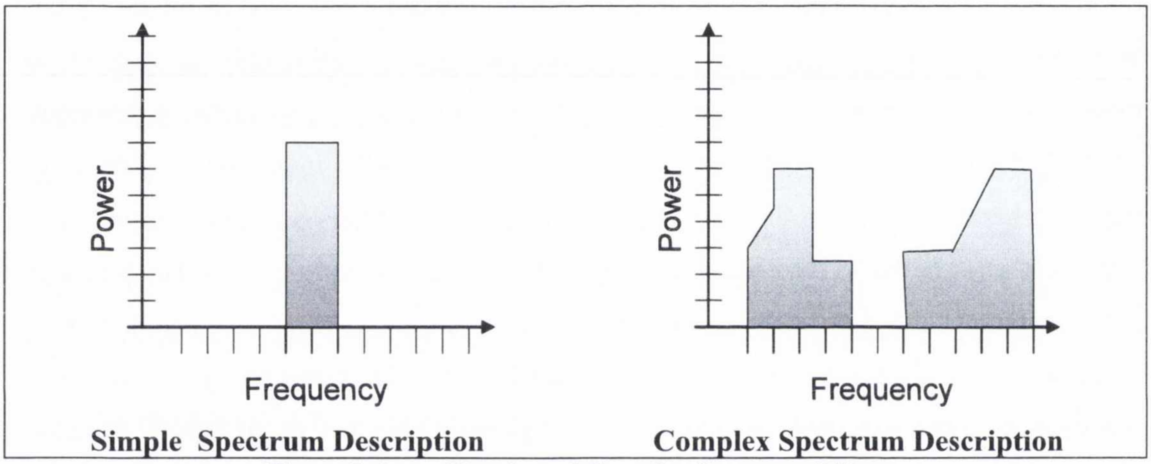

Simple Spectrum Description **Complex Spectrum Description**

Figure 2-6 Simple and Complex Spectrum Objects

Similarly, the *Waveform* object can be as simple as a description of an FM waveform with just center frequency, frequency deviation, and power. Alternatively, it could be a complex description of multi-carrier, multi-antenna exploiting waveform, along with harmonic and intermodulation product energy masks. In the limit, it could consist of a hybrid wideband signaling waveform, such as contemplated by Wu and Natarajan [93], [94], Chakravarthy [95], or a mix of techniques, such as proposed by Rajbanshi et al. [96]. This complexity demonstrates need for the extensibility of class or ontology structures, since the number of variations is seemingly infinite, and many of these require unique semantics to fully describe the characteristics.

As an example, the following section outlines the decision process that a cognitive radio might execute in determining the most beneficial physical layer choices. This discussion is not intended to be exhaustive, but is to provide an anecdotal framework for how the results of the quantifications provided in the upcoming chapters fit into a network-focused decision process.

The process of determining permissible spectrum options will yield a set of spectrum options, such as (simplistically) shown in Table 2-5 (presumably by some exogenous interference avoidance policies).

In addition to the permitted frequencies of operation, it is proposed that the cognitive radio algorithms also use the spectrum sensing information to provide an estimate of the in-channel noise in each frequency option, and an estimate of the likely intermodulation noise resulting from mixing in the front-end of the receiver. Methods for the determination of the intermodulation energy will be discussed in detail in Chapter 4.

Table 2-5 Example Spectrum Options

Frequency		**Energy Limits through Band**[7]			**Estimated Noise Floor**	
From	**To**	**Energy (From)**	**Energy (To)**	**Max Power Emitted**	**In Channel**	**Overload**
(MHz)	**(MHz)**	**dBm/Hz**	**dBm/Hz**	**dBm**	**dBm/Hz**	**dBm/Hz**
1410	1450	-50 dBm	-40 dBm	30	-172	-168
1450	1780	-40 dBm	-40 dBm	30	-172	-220
2000	2150	-35 dBm	-35 dBm	20	-168	-280 dBm
2150	2180	-35 dBm	-30 dBm	20	-166	-150 dBm

The first frequency range (1410-1450 MHz) in this list offers a very low noise floor, but a prediction that adjacent channel intermodulation may raise the noise floor by approximately 5 dB. The maximum energy in the lower end of this band must be reduced, presumably to avoid adjacent channel interference to adjoining spectrum users. The second range (1450-1780 MHz) has the same low noise floor, with a predicted intermodulation level that is well below the receiver noise floor, and thus, this noise can be discounted. In this range, a constant emitted power level is permitted. The third frequency (2000-2150 MHz) has considerably increased in-channel energy (a 4 dB higher noise floor), but very low level of predicted front-end intermodulation. A constant energy density is permitted in this band. The last range (2150-2180 MHz) has slightly increased in-channel noise (but still presumably policy compliant, as it is only 6 dB above the noise floor, here assumed to be in the region of -172 dBm/Hz), but has a probability that the receiver will encounter an additional 22 dBm of intermodulation energy due to the presence of adjacent channel energy.

The cognitive radio analyzes each spectrum opportunity to determine which of its operating modes is capable of operating within its constraints of bandwidth and energy. For a given frequency candidate, some modes may fail to provide link margin, and some

---

[7] Two energy density values are provided in order that the table express a ramp of permitted emission level across a single band, such as provided by the US UWB emission mask.

may have no waveform modes that do not exceed spectrum thresholds for operation. Table 2-6 illustrates this structure, using a radio with a small choice of waveform options, including Binary Phase Shift Keying (BPSK) and two modes of Quadrature Amplitude Modulation (QAM) as examples of the process.

Table 2-6 Example Waveform Options

Frequency		**Waveform**			
From (MHz)	**To (MHz)**	**Modulation**	**Bit Depth Bits/Hertz**	**Width MHz**	**Energy (dBm)**
1410	1450	BPSK	1	40	20
1450	1780	BPSK	1	10	30
1450	1780	QAM	2	10	30
1450	1780	QAM	6	10	30
2000	2150	BPSK	1	2	25.6
2000	2150	QAM	2	3.2	30
2000	2150	QAM	6	3.2	30
2150	2180	BPSK	1	0.5	20
2150	2180	QAM	2	1	20
2150	2180	QAM	6	1	20

The frequency option starting at 1410 MHz is so constrained that only a minimal amount of output power is permitted. Although permitted, this band choice is highly undesirable, even at the highest end of the band. Presumably, the OFDM mode of operation may not be viable due to the reduced power permitted, and the required linear mode of operation. The 1450-1780 MHz range has a number of viable waveform candidates that can utilize the full power permitted, with wide modulation range and symbol throughput. In the 2000-2150 MHz range, the BPSK mode is slightly constrained, and in the 2150-2180 MHz range, the limitations at the upper end of the band constrains the total power. The topology layer considers the effect of the frequency choices and waveform options on the (predicted) network reachability and throughput. A major criterion is the reachability of the nodes, given each frequency and waveform option. The objective of the network is not necessarily to maximize connectivity; it may be that lower reachability allows for increased inter-node bandwidth, a more stable topology, and reduced network complexity. These are not design choices, but ones that are inherently situational. Chapters 6 and 7 will demonstrate that dense network performance can be enhanced with shorter, higher throughput links.

An example topology option set is provided in Table 2-7. Note that in the absence of awareness of the mission of the radio, it is not possible to decide which of these options provides the best strategy for a given situation. While the focus of this dissertation is awareness of the physical environment, the equivalent network processes equally require cognizance of user or network objectives in order to balance and assess these alternatives. The alternatives are to maximize connectivity, or maximize throughput, or seek a hybrid of each. Approaches for determining optimal selections of modulation and coding are provided in work such as by Cui, Goldsmith, and Bahai [97].

Table 2-7 Example Topology Options

Frequency		**Waveform**				**Node Reachability**
From (MHz)	**To (MHz)**	**Modulation**	**Bit Depth Bits/Hertz**	**Width MHz**	**Energy (dBm)**	
1410	1450	BPSK	1	40	20	#2, # 3
1450	1780	BPSK	1	10	30	#2, # 3, #9, #15, #20
1450	1780	QAM	2	10	30	#2, # 3, # 9, #15
1450	1780	QAM	6	10	30	#2, # 3, #9
2000	2150	BPSK	1	2	25.6	#2, # 3, #9, #15
2000	2150	QAM	2	3.2	30	#2, # 3, #9
2000	2150	QAM	6	3.2	30	#2, # 3
2150	2180	BPSK	1	0.1	20	#2, #3
2150	2180	QAM	2	1	20	#2, #3
2150	2180	QAM	6	1	20	#2

In this example, the first frequency and waveform option (1410-1450 MHz/BPSK) has very limited reachability due to the low power levels permitted at the band edges. In the second range option (1450-1780 MHz), the low order, constant envelope modulation (BPSK) offers the greatest reachability, and the QAM modulation offers more bandwidth, although, at reduced reachability due to the higher signal to noise ratio needed for the higher order constellations. Also, note that the high intermodulation induced noise floor, shown in Table 2-5 significantly compromises reachability for all modes in the 2150-2180 MHz band. The important dimension that cognitive radio and network technology brings is the recognition that decision between these modes can only be made with awareness of the nodes and network's operating needs. For example, if there is a large amount of traffic to node #20, then the 1410-1450 MHz BPSK option has advantages in minimizing the number of hops needed to deliver the traffic, although at a cost in bandwidth available. If, on the other hand, most of the traffic is with nodes #2 and #3, then maximizing the throughput, and accepting the need for routing to nodes #15 and #20 may be the best use of resources. Quantifying these constraints and assessing the specific performance impacts is the subject of Chapters 4 through 9.

This structure fully implements the partition between endogenous and exogenous processes. The regulatory community will have an interest in the spectral usage of a cognitive radio in order to assure non-interference to the other users of the spectrum. Spectrum policy enforcement would validate the *Spectrum* object contents as first criteria for further consideration by the cognitive radio. The cognitive radio endogenous optimization can determine waveform and reachability implications in isolation of the regulatory concerns. Isolation of the interference avoidance policy from the radio algorithms that determine the best utilization of the permitted emission is important to enable cognitive radio technology to rapidly evolve without continuous regulatory involvement. This structure is a compatible generalization of the structure proposed by the author [98], [99] and reported [74], [75] in the DARPA XG DSA technology experiments.

This cognitive radio process requires that the relative desirability of each emission option be continuously assessed to determine its desirability relative to other available choices by application of some metric of desirability. A partitioning of this ranking metric has four fundamental knowledge domains: the decision option; the known characteristics of the spectrum and network environment; the devices design (capabilities, limits, resources, …); and the desired mission.

Metric = f (<evidence>, <capability>, <emission>, <mission>)

Where:

<evidence>	The collection of sensory information (and potentially history) characterizing the environment of the cognitive radio. Some of this evidence may have also been used to obtain permission to radiate.
<capability>	The characteristics of the radio and its constraints on the operation, or the desirability of operating modes. These could include power amplifier back-off, filter bandwidth, LNA linearity and other factors that the radio must consider in weighing strategies, or that act as constraints on the operating envelope.
<emission>	The emission tuple described previously (spectrum, waveform and topology). A separate process may validate the permission to radiate this energy, as provided by the PCR discussed previously.
<mission>	The characteristics of the offered load, the operation of the link, the time criticality of the information and other service level constraints that the radio must consider in optimizing its performance. For example, retransmissions may be efficient for file transfer, but not as effective for streaming transfers, such as voice and video.

The next section describes an algorithmic model for progressive resolution of the appropriate operating point and the selections of adaptations. Each functional decision of the cognitive radio selects from sets of discrete filter, frequency, mode, power, and other characteristics of operation. These adaptations are generally well known elements in the design of wireless devices, and are traditionally static design choices. The cognitive radio algorithms provide the processes to dynamically select specific adaptations based on real-time assessment of mission needs and environment. In effect, these adaptations defer design decisions until the moment of execution, when that can reflect the actual environment of the device. The process is similar to manual design; the difference is in its time-scale, and its environmentally aware implementation.

It is possible that this decision process could be defined as a linear, stateless optimizing engine, considering all possible permutations, or as an inference engine that performs constraint solving. However, for discussion here, the process is modeled as incremental decision and sensing operations. This process reflects that the operation of a cognitive radio is also a Decision Theory problem. The device is constantly offered the opportunity to expend resources (energy, channel time, processing cycles) to acquire more information about its environment, and acquire additional confidence in the environmental assumptions. The algorithms must determine which environmental awareness sensing opportunities[8] actually have value in achieving an improved operating point selection. In addition to the decisions commonly ascribed to a cognitive radio, it must also make decisions regarding the incremental value of additional environmental sensing based on a state of environmental awareness that can range from ignorance, statistical estimates, to high confidence (the value was known recently, and it is unlikely to have changed). This process is reflected in the structure to follow.

For example, if the device has the capability of exploiting multiple path propagation, as is provided by Multiple Input/Multiple Output (MIMO) [100], the device must understand the multipath conditions (such as signal delay spread) in order to determine the benefits of committing the resources to operate in a MIMO mode. Although a first order view would consider this a single decision, it really is two dependent decision layers, as depicted in Figure 2-7. The first decision is whether to perform the sensing at all. This decision has several outcomes, driven by the devices initial appraisal of its own knowledge of the channel. It is reasonable to assume that if recent multi-path data had been sensed, the device would have high confidence in its knowledge, but that this confidence would decay over time. This confidence could then justify making a decision without spending any resources on the (likely redundant) sensing operation. If, on the other hand, it had no knowledge of the channel, it would first have to determine the benefits of the MIMO operating mode before deciding to invest the

---

[8] Most cognitive radio descriptions consider sensing as focused on the spectrum or the signaling channel. In the context used in this dissertation, it also includes sensing of the network and application environments, traffic and performance needs.

cost of further refining its knowledge, or creating a high confidence channel estimate.

Deciding the marginal utility of multipath exploitation is an extreme case of the Decision Theory that is implicit in the operation of a cognitive radio; but the same situation is present in cases of sensing an additional spectrum band, or selecting resolution bandwidth when conducting energy sensing. Efficient cognitive radios will require more research to ensure they do not fall victim to "analysis paralysis" by over expending resources for sensing.

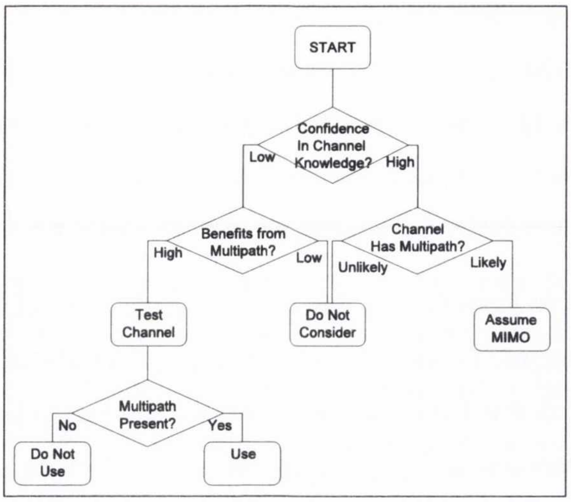

Figure 2-7 Typical Decision Process for Decision to Expend Sensing Resources

A meta-structure for this processing is shown in Figure 2-8. In this example, a single policy conformance reasoner is utilized for spectrum policies, but as shown previously in Figure 2-3, there could be multiple of these to reflect operating limitations imposed by multiple domains.

The figure illustrates that the environmental sensing at each step is contingent on the previous decisions. The device only invests energy, channel time or other resources to determine the viability or desirability of attractive candidate actions. This reflects the philosophy described in Figure 2-7. The order of the decisions should reflect the relative number of candidates that are excluded by each sensing/decision step, and the relative resource cost of the sensing and decision process. In the chapters that follow, these decisions are arranged in a logical flow, but the actual sequence would be optimized for the specific resource usage of a given implementation. The model of the process as successive tuple filtering implies that the process

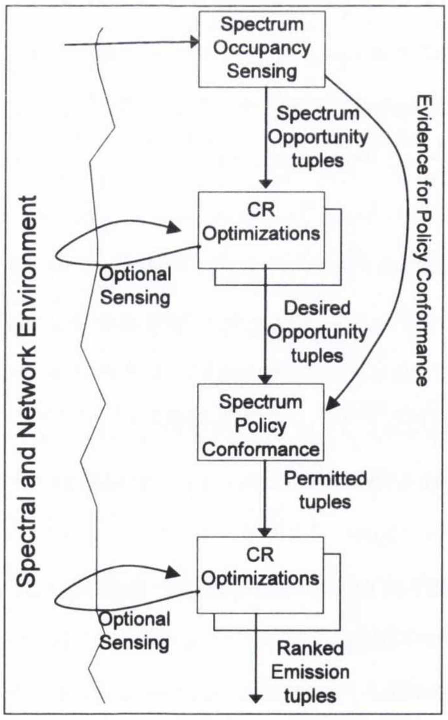

Figure 2-8 Meta-Structure for Cognitive Radio Processing

is composable, and that the actual sequence of operations is not critical to recognizing optimal solutions, so long as solutions sets are not prematurely truncated. It is an issue for the implementation to ensure that it can implement a metric that assesses the distance of a given solution from a potentially optimal one before truncating the search.

The Spectrum Occupancy sensing serves to provide the initial candidate emission tuples, and also the evidence by which the policy conformance checker will ensure compliance of regulatory constraints. This depiction assumes the regulatory compliance is a separate logical function from cognitive radio optimization. Although they both operate on sensing data, their intent is fundamentally different, and perhaps conflicting so it is not unreasonable to assume that they are not an integrated mechanism. The conformance reasoners represent the exogenous constraints, and the optimizers represent the endogenous interests of the device.

For generality, some of the optimizing processes are shown as being performed before the emission policy conformance is obtained, and others after conformance testing. There is no fundamental difference in these approaches; the decision regarding the placement of each function is assumed to be based on the relative resources required to perform emission policy conformance testing and of the cognitive radio optimizing processes. All possible emissions could be checked for policy conformance before any optimization processes were performed, or alternatively, only the selected emission could be checked after all other processes were completed.

It is also important to recognize that although the implicit assumption in this process is that a single frequency opportunity is the target of the search, it is possible that more than one physical or logical channel will be desired. A discontiguous waveform (such as OFDM with subcarrier removal or multiple QAM carriers) is one example of this possibility, but a more interesting case is a cognitive radio with multiple transceivers. The author initiated the WNaN program at DARPA with one of its objectives being to transition to cognitive radios that simultaneously managed a number of transceivers with diverse channel operating characteristics and network topology roles [10]. In this case, a number of emission tuples are needed to satisfy the requirements of all transceiver links. In this situation, the tuples are assessed with a number of different parameters in <***mission***>, reflecting different roles for each transceiver within the network topology. For example, one might be focused for local communications within a cluster of nodes, while another is providing more long range, lower data rate backhaul to an infrastructure.

2.5 Algorithmic Description of Cognitive Radio Processing

The preceding discussion provided a very general model of the analysis of the essentially infinite options for radio operation, and the process to select a subset of one or more that maximized the mission utility function of the device. This section provides an initial partitioning of these processes; and then provides the structure for the discussion of specific cognitive radio methods and performance in Chapters 4 through 9.

This discussion follows the general framework of Figure 2-8, and replaces the abstract algorithm representations with the specific methods that comprise the logic of a cognitive radio, maintaining the interactive sensing process that was depicted in Figure 2-8. The depicted sequence of operations reflects the principle (established in Section 2.4) that the process must balance the sensing resource cost, with the degree to which the information constrains the possible solution set, when determining the sequence of operations to be conducted. Figure 2-9 illustrates a possible logical processing structure for a cognitive radio.

An important common characteristic of each block is that they each represent a standalone function that can be assessed and quantitatively characterized in terms of its contribution to either increased cognitive radio capability, or reduced requirements for hardware component level performance, such as front-end linearity, battery energy storage, or power amplifier capability or linearity.

The sequence of this figure is not arbitrary, but is derived from a notional estimate of the relative cost of the sensing operations involved. Band energy sensing is the lowest cost. It has low-resolution bandwidth, is not sensitive to the noise floor, and in all probability, may disqualify the largest extent of spectrum. High resolution, noise floor sensing is performed next. This requires much higher resolution bandwidth, and operates close to the natural noise floor. Inherently, this is a much more expensive operation. Lastly, channel specific sensing is performed. Unlike the other sensing operations, this cannot be performed unilaterally, and requires that at least a pair of transceivers be set to the frequency, and initiate communications. Clearly, this is a resource consumptive sensing operation for a cognitive radio. Note that the sensing for information that is internal to the wireless device is not depicted, as it is assumed that the cost of obtaining this information is not significant in the total resource usage, and therefore not a consideration in the organization of the process.

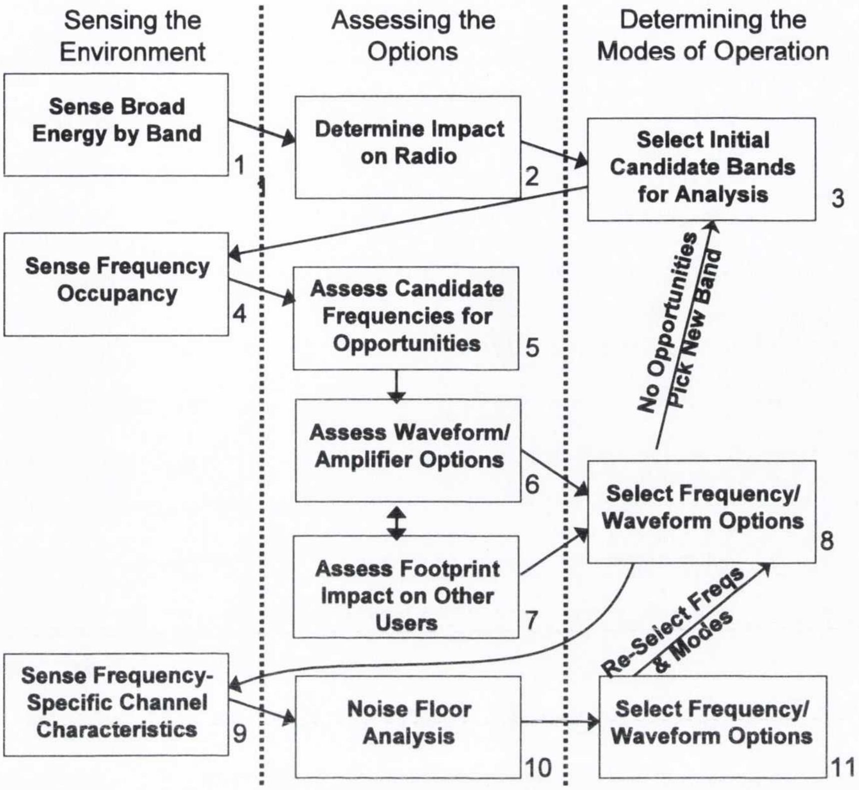

Figure 2-9 One Possible Cognitive Radio Physical Layer Processing Model

The following discussion proceeds through the process depicted above, referring to each numbered function. It is worth recognizing that this process is one that is composable in general, and different sequences of operation are certainly possible, but do not change the character of the individual processes.

The first operation the radio will perform is an analysis of the aggregate energy contained in the available or candidate pre-selector bands to determine the total energy present in each (1). Previous work has generally considered DSA as a technology to assign frequencies for radio operation in a manner that avoided co-channel interference. This dissertation extends this concept to also include consideration and mitigation of adjacent channel effects. (The physics of this will be discussed in detail in Chapter 4). These effects arise because the total energy that is transferred from the antenna exceeds the linear range of the receiver's front-end amplifiers. When this occurs, all of the signal products intermix, creating second and third order mixing products. In particular, the third order products occur within the filter pass-band frequencies, and for complex spectral environments, create a large amount of intermodulation noise that raises the effective noise floor[9]. Although a frequency may be unoccupied, and thus suitable for selection

---

[9] Note the engineering test for this is a "two tone" test that produces tonal intermodulation products of only two signals. However, in a complex environment with a large number of signals with modulation present,

based on both channel characteristics and spectrum policy, the total induced intermodulation energy may have significant detrimental impact on its suitability. In this dissertation, the spectrum management objectives of DSA are extended to include the sensing and decisions required to limit the amount of energy that the front-end is exposed to. The sensing operation will consider the total energy that would be received with each pre-selector filter setting option, instead of looking at individual frequencies.

Marshall postulated [101] that the benefits from cognitive avoidance of front-end overload could be as significant as the spectrum access benefits arising from the dynamic frequency management provided by DSA. The sensing operation required to support band selection can be very low resolution, since its intent is to assess only the total energy that the Low Noise Amplifier (LNA) is exposed to in each potential band of operation[1]. These options will be compared to the known performance characteristics of the front-end (2) to determine the range of intermodulation effects that would occur in each pre-selector band option. The likely intermodulation level would be determined, and the additional noise energy estimated.

This process culminates (3) in a quantified assessment of front-end overload and overload effects (additional broadband noise), enabling the radio to develop a list of frequency bands that are suitable for operation, and a quantification of the noise floor escalation of each possible option. The selection process can be more complex than just selecting the lowest energy bands; it can consist of a more subtle process that assesses compromises between the effects of additional noise or desensitization[2] due to strong signals with other desirable considerations, such as propagation and antenna performance. As an example, it might be an attractive option to select a frequency an octave lower than another option, even having to accept an additional 2 dB of intermodulation noise to increase the capture area of the antenna by 6 dB[3].

---

the intermodulation process creates mixing products that approach a more normal Additive White Gaussian Noise (AWGN). In this dissertation, the most concern is the third order products passed through a less than octave-wide filter, some of which map back to the original band of interest. For simplicity, the mixing products are modeled as AWGN, although some degree of spectral or temporal correlation may be present. In fact, this correlation may increase the total intermodulation energy, as will be discussed in Chapter 3.

[1] This decision is assumed to be local to an individual node, since the front-end energy is often driven by proximity to other emitters, rather than the general character of the spectrum.

[2] AGC desensitization occurs when the receiver senses the total energy in one if its receiver stages, and reduces the gain to maintain linearity. The effect of this is to avoid intermodulation, but also to reduce the effective sensitivity, and thus increasing the effective noise floor. This process is effective and of no significant consequence if the energy is from the intended signal, but is detrimental to operation if the energy is from other signals, and the reception of the intended signal is compromised.

[3] This 6 dB is the theoretical increase in effective antenna aperture for an isotropic antenna when halving the resonant frequency.

Note, this dynamic management of dynamic range is a unique capability that only a cognitive radio can implement. In conventional radio operation, the frequencies are generally assigned in advance. Therefore, in an uncorrelated environment, there is a random chance that the assigned frequency will fall in the same pre-selector band as a strong single signal, or cumulative energy, that will drive the LNA into a region where it will generate strong intermodulation products. When DSA is utilized, the probability of LNA overload can be reduced with the additional complexity of also assuring that adjacent energy will not introduce unacceptable intermodulation. This is an important, and less recognized advantage of cognitive radio for two reasons. First, since even the highest performance radios cannot assure non-overloaded operation when assigned to arbitrary frequencies, it enables a cognitive radio to provide essentially guaranteed linear receiver front-end operation. Secondly, it can significantly reduce the intrinsic front-end linearity required to provide at least equivalent probability of overload. Although communications analysis generally stresses efficient use of transmitter power, in many applications, the energy consumption of the high performance receiver components are the dominant user of energy, and this energy use is proportional to the linearity required.

Specific frequencies are selected based on a high resolution and sensitivity scan of candidate bands of operation (4). While the resolution of the earlier pre-selector band scanning needed only to be sufficient to identify total energy in each band of interest, the resolution of the frequency selection scan is driven by the likely occupied bandwidth (for both incumbent and generated signals), and the spectral noise sensitivity that must be achieved.

The actual assessment of frequencies has a number of considerations (5). In a cognitive radio, one of the primary functions is to determine which frequency is appropriate for use. Although this dissertation is not addressing the regulatory issues associated with cognitive radio, it is apparent that a well-behaved protocol would desire (or be compelled) to avoid interference to other users, as well as avoid co-channel interference on frequencies that are under consideration for operation. In a complex implementation (as will be assumed in Chapters 6 and 7), this list is a set of allowable power spectral densities that could be utilized by the radio, in one or more modes of operation, as shown in Figure 2-6. While the initial consideration of DSA systems has generally assumed that the waveform was fixed and known, and the radio was searching for "white space" in which to locate the signal; this dissertation considers a more general case in which the cognitive radio considers the waveform selection to be a consequence of the spectrum opportunities it is provided.

Although the most suitable frequency might be an obvious choice (and the one selected in some implementations), there are reasons that argue for more nuanced ranking of the candidate frequencies. Two critical aspects of this decision are shown. One aspect is waveform selection, depicted in (6). This consideration will require matching the information carrying capability of the waveform to the available spectrum opportunities; the radios operating limitations (such as peak power, battery energy, etc.); any known channel conditions, such as multipath, that would drive the selection of the waveform; and any constraints arising from the performance of the device in handling the waveform selection (examples include phase stability, peak and average thermal output, linearity, etc.). For example, the radio could trade between the improved Inter-Symbol Interference (ISI) of an OFDM waveform (multiple tone waveform with very low symbol rate) against the effect of significant back-off (typically in the range of 8 dB) in amplifier output to accommodate the waveform's PAR. Both choices have their advantages, and the only way to select between these strategies is to assess the specific situation and its needs and challenges. This is the advantage of cognitive radio.

Another consideration in waveform selection is the effect of the waveform footprint on the other users of the same band (7). This is important if the radio's network needs to enable a large number of devices to operate cooperatively or independently in a high-density environment [102][1]. If each radio network is independent, then the operator may not elect to consider the effect on other users. When the spectrum environment contains a single operational interest or entity, such as military users, public safety institution(s), or a coordinated activity, such as organizations participating in a natural disaster recovery, then the operator of a set of cognitive radios will acquire an interest in not just optimizing the performance of individual users; it will have an even greater interest in optimizing the performance of the aggregate set of radios, even if this mode is less than optimal from the perspective of a single radio. With this added context, the cognitive radio algorithms can

---

[1] A number of game theoretic models have shown that the assumed "Tragedy of the Commons" is not an inevitable consequence of unregulated access to the spectrum. This dissertation will not address the reasons why a radio would desire to consider its effect on other users, since these are organizational, economic, and regulatory in their nature. This dissertation will address the means for implementing this consideration. This is therefore not an unlikely behavior that certain policies might desire to implement. Alternatively, an industry standard could force compliance with spectrum etiquettes as a part of interoperability testing, such as by industry consortiums.

make some initial decisions regarding its operational modes, or at least further refine the set of candidate modes for later selection (8).

It is possible that all of the operations described above will lead to a number of operational alternatives (frequency, waveforms, amplifier modes, etc.) that are operationally suitable, and not strongly differentiated. In this resource rich case, additional considerations may serve as "tie breakers" to provide additional benefits from cognitive radio operation. At this point, the sensing of the radio has to transition to more active, and resource consuming mechanisms to obtain more subtle characteristics of the channel (9). The simpler ones (and lower resource cost) to measure include potential sources of link interference, such as impulsive noise, spur signals, and elevated noise levels. These attributes of the channel can be detected by the nodes in a unilateral mode; it requires no cooperation from other nodes in the network. Potential channel impacts are assessed and weighted (10).

At this point in cognitive radio planning, the decision process has; successfully precluded interference to other spectrum users, found suitable frequencies, selected among these based on the capabilities of the radio, further examined the individual characteristics of the most optimal channels, and is now in position to make the optimum decision for radio's operation for the next interval of time (11).

This discussion implicitly treats the decision process as within the unilateral control of a single node. In reality, it is at least a bilateral decision between link pairs, and most likely multilateral, with multiple members of the network. Previous work by Fischer et al. [103] has demonstrated that an arbitrary number of nodes that are operating through a broker process reach convergence towards a decision that is within the acceptable decision criteria of all participants, and that this algorithm is scalable, as convergence time increases only logarithmically with the number of nodes.[1]

2.6 Towards a Unified Theory of Interference

A key objective of DSA has been to demonstrate that DSA mechanisms can avoid interference to other users of spectrum bands in which they could be permitted to share on a non-interference basis. This is an important objective, since during the initial deployment of cognitive radio, it is likely that the number of cognitive radios will be

---

[1] As will be shown in Chapters 6 and 7, the number of devices that need to be physically rendezvoused onto a frequency does not need to grow with the size of the network, and in fact is best held constant through use of connected sub-networks.

small, and they will be a minority in any band in which they operate. This first generation of cognitive radio will have to maintain the key principle of current spectrum management practice in protecting users of the spectrum from other users.

However, when looking beyond the initial cognitive radio deployment, it is possible that cognitive devices, with interference mitigating technologies will become a majority of the device population, particularly if the economic arguments for their deployment are realized in practice. The interference mitigating characteristics of these devices could then be employed more aggressively; providing spectrum self-relocation in response to interference, in consonance with other interference avoidance and channel management functions.

The wireless communications community addresses interference from a number of perspectives, and in a variety of domains. These include manual spectral separation of signals, time-domain sharing of spectrum (such as listening for a clear channel, or fixed slots), DSA, MAC protocols, such as Carrier Sense, Request and Clear to Send (RTS/CTS), and more recently MIMO[1]. Another structure for considering these mechanisms is to categorize them in relationship to the other spectrum users from which they isolate the radio link, as shown below.

Other non-Cooperative Users	Provides absolute separation of spectrum usage in time or frequency. Provides as close to orthogonal operation as technically and procedurally possible, at least for primary users.
Other Nodes of the Same Network	Media Access deconfliction allows cooperation amongst users of a single channel, such as in Ethernet-like WLAN protocols, trunk servers, cellular systems and similar designs that assume cooperation among users sharing a channel.
Multiple Antennas on the Same Node	MIMO technology provides separation of signals from multiple antennas on a single node.

In current practice, each of these operates with very different risk tolerances and throughput models; even though they each contribute to measures of communications effectiveness that are scalar in many senses. Since reliability and performance can be traded off in communications systems, it stands to reason that each layer of the interference avoidance process optimizing for independent standards of interference

---

[1] The use of Direct Sequence Spread Spectrum of various forms is an even more general formulation of this structure. However, it is only possible when codes from all users maintain their orthogonal character when they arrive at all other users, and when the range ratio of near and far terminals can be controlled. It is therefore a highly useful construct, but not one that can be the basis of general rules of operation.

avoidance is sub-optimal. Although there are few quantitative criteria for interference avoidance through manual spectrum planning, it is intended to cause interference events to be extremely rare, so an estimate of 10^{-4} is not unreasonable. By contrast, MAC layers typically operate at collision rates on the order of 10^{-2} to 10^{-1}. Some proposed approaches to DSA in fact blend concepts of non-cooperative and cooperative nodes, and mix DSA and MAC layer approaches, such as by Marshall [6], [87], [104]. This discussion considers MAC layer techniques as focused on nodes with overt cooperation, typically with a significant degree of protocol commonality, and mutually readable traffic. DSA, in contrast, provides some degree of sharing, but without the constraints of common protocols, or any degree of information sharing, or even awareness of the other users.

An optimal solution should consider all forms of interference within an integrated framework that could assess the performance benefits of allowing more interference from one cause, versus allowing it from another, or not allowing interference at all. Although individual nodes operate best in a non-interfering environment; a collection of nodes may well achieve its optimal aggregate capability when local levels of interference are allowed to achieve levels higher than those associated with current spectrum practice. Chapter 6 will illustrate this principle in terms of interference between DSA capable radios, but the principles can, and should be extended to include all layers of the interference avoidance process. Figure 2-10 illustrates this progression.

An example of this is recent work that analyzed the use of MIMO processing to reduce interference from emitters, instead of being utilized to increase the total throughout of a single node-to-node channel set. Ozgur et al. [23] showed that theoretically the throughput of an ad-hoc network could be essentially constant, so long as perfect awareness of the channel can be achieved with additional antennas. This result is comforting in that it provides at least a theoretical argument that the limitation described by Gupta and Kumar can be overcome by suitable quantities of antennas. Unfortunately, the proof does not lead directly to an engineering solution. Similarly, Govindasamy et al. [105] report that significant benefits in interference avoidance and total throughput of a

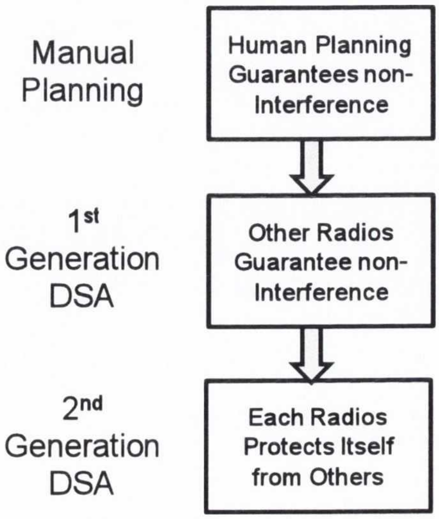

Figure 2-10 DSA Evolution

set of nodes can be achieved through use of multiple receive antennas, even in the absence of channel state information. This discussion partitions the universe of spectrum dependent systems into two categories:

Interference Intolerant	Systems that cannot internally mitigate interference events within a range of resources and/or operational disruption that is acceptable to the system and its uses.
Interference Tolerant	Systems that can tolerate frequency and time collisions, and can mitigate them, potentially at some cost to the device, but within a range of costs that are acceptable within their mission constraints.

An example of how such an integrated process might operate is provided in Table 2-8. Each interference mitigation technique has an associated cost, to perform the mitigation process, and through the impact of events that are not successfully mitigated. Interference events cost not only the loss of individual communications events, but also the system actions to address the interfering conditions.

Table 2-8 Functional Element Participation in Integrated Interference Avoidance

Functional Element	Responsibility within the Unified Interference Concept
Manual Spectrum Planning	Provide isolation of interference tolerant and non-interference tolerant systems so devices can make assumptions about the degree of interference that they are permitted to cause to other users of the spectrum. Instead of "protecting" individual users, the spectrum planning process protects classes of users.
Dynamic Spectrum Access	The DSA functionality determines the probability that a given frequency selection will receive interference events, or cause them to others. Depending on the frequency, the tolerance to causing interference is driven by the policy established through the manual spectrum planning process, and the tolerance to receiving it is driven by the minimal Quality of Service and the other alternatives available to the device.
MIMO (if Present)	An alternative to time and frequency separation; exploits spatial dimensions to isolate signals in order to receive multiple simultaneous information streams, or to isolate a signal in the presence of interference that would preclude demodulation.
Media Access Layer	Provide temporal separation of signals in a common signaling channel, through shared protocols or channel access procedures and allocations.

A summary of the cost to maintain interference avoidance and the cost of failure of these mechanisms is provided in Table 2-9.

Table 2-9 Interference Avoidance Maintenance and Failure Costs

Functional Element	Cost of Maintaining Interference Avoidance	Cost of Interference Avoidance Violations
Manual Spectrum Planning	Extensive manual planning and suboptimal and static allocation of spectrum, with a necessity to provide high degree of protection for non-interference tolerant systems.	Complete failure of non-interference tolerant systems. Minor impact to interference tolerant systems.
Dynamic Spectrum Access	Periodic spectrum sensing, consuming some channel time, and additional signal processing.	Short-term loss of a single packet, longer-term interference forces frequency relocation of all members of a network.
MIMO (if Present)	Requires commitment of multiple, independent transceiver channels to separate multiple streams, as well as additional processing to perform signal processing.	Loss of individual packets, so long as the channel conditions are sufficiently dynamic to create Eigen-separation of channels. Otherwise, the level of MIMO processing must be reduced.
Media Access Layer	Additional message traffic and/or lost channel access opportunities to schedule use of the channel through fixed (slotted) or random (RTS/CTS) access.	Loss of multiple transmissions (both colliding frames), and required repeated attempts.

A simple example of integration of Physical layer and MAC layer cooperation could be an operating mode that provided "*in-cast*" services[1], in which multiple nodes could simultaneously transmit information to a single node, such as in a sensor fusion scenario, with the MAC scheduling the transmission, and MIMO separating the signals.

This probabilistic treatment of interference is a logical extension of the structure shown in Table 2-5 through Table 2-7. Table 2-5 can be extended to reflect the probability of both being interfered with, and causing interference, as shown below.

Table 2-10 Probabilistic Extension of Spectrum Availability Table

Frequency		Interference Probability	
From (MHz)	To (MHz)	To Others	To Self
1410	1450	10^{-4}	10^{-5}
1450	1780	$10^{-2.5}$	10^{-3}
2000	2150	10^{-2}	$10^{-1.5}$
2150	2180	10^{-3}	10^{-4}

---

[1] In-cast is considered the opposite of broadcast. Broadcast is considered as one transmission being received by multiple nodes; in-cast is multiple transmissions from different nodes being received by one node. Unlike MIMO, it is assumed that the transmissions are from separate nodes.

It might be assumed that the first option (1410 MHz) was a band that was shared with interference tolerant, or mitigating devices, while the other bands were shared with interference in-tolerant devices; although the cognitive radio logic does not need to concern itself with these policy matters once the exogenous reasoning has determined the availability of the opportunity. This formulation supports algorithms that are both sensitive to the mean value of anticipated performance, as well as its probability distribution. It inherently reflects both reliability and performance, and the compromises between them.

To complete a model of cognitive radio performance, it is necessary to develop a theory of radio operation that integrates the effects of decisions across layers, rather than abstracts them, as in the conventional network model. A theory that is inclusive of multiple channels, multiple antennas, methods of access control, methods of coding, and dynamic topologies is critical to understand the fundamental limits [106] of cognitive radio.

2.7 Quantification of Cognitive Radio Contribution

To assess the impact and benefits of cognitive radio, it is necessary to aggregate the impacts of adaptive operating decisions on the anticipated performance of the cognitive radio compared to the performance of implicit or default decisions of the equivalent non-cognitive radio. This provides a quantification of the improvement in the mean performance of devices with similar hardware capability, and the impact on the confidence of achieving specific performance levels (reliability). For each of the parameters in Table 2-4, the relations below provide an estimate of the technology's contribution.

$$\text{Performance Improvement} \approx \frac{\text{Mean Cognitive Radio Performance}}{\text{Mean Conventional Radio Performance}}$$

$$\text{Reliability Improvement} \approx \frac{(1 - \text{Conventional Radio Probability of Performance})}{(1 - \text{Cognitive Probability of Performance})}$$

Figure 2-11 provides a representative performance curve for the three radio devices that will be used to describe the methodology of determining reliability. *Performance Optimized Cognitive Radio1* and the *Non-Cognitive Radio* have the same underlying component capability, so the difference in reliability and performance between them is strictly due to the effect of the cognitive adaptations. *Cost Optimized Cognitive Radio* has one or more of its component performance measures reduced to achieve the same minimum performance as the Non-Cognitive Radio. The reliability of each radio is reflected by its intersection with the minimal acceptable performance line. The improvement in reliability for equivalent component performance is measured as the difference in reliability achieved by *Performance Optimized Cognitive Radio* and the *Non-Cognitive Radio*, points *A* and *B* respectively. As will be shown in the later chapters, the major improvement achieved through the adaptive algorithms is in the low performance/highly stressing regions where adaptation is critical to achieving even minimal performance levels. An alternative exploitation of cognitive radio technology is represented by the lower component performance. *Cost Optimized Cognitive Radio* has one or more of its component performance measures reduced to achieve the *2*. This radio uses cognitive adaptation to achieve equivalent reliability (point *B*) at lower component performance levels. A depiction of performance impact determination is in Figure 2-12.

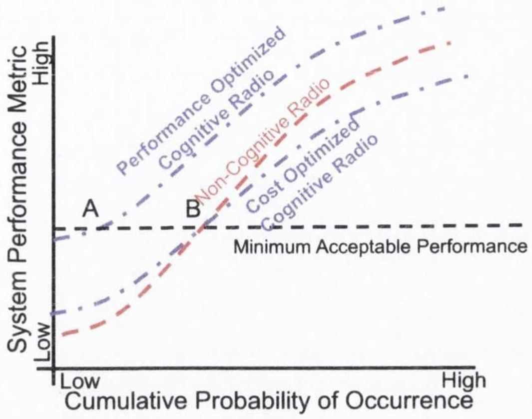

Figure 2-11 Representative Mode of Operation –Minimal and Optimal Performance Points

The line depicted in *A* represents the performance increase that is provided through the inclusion of cognitive radio features. This line provides the change in radio performance for a constant level of component performance. Line *B* provides the decrease in component requirements necessary to achieve a given level of radio performance. Applied together, the approach shown in Figure 2-11 and Figure 2-12 provide the framework needed to determine the performance matrix shown in Table 1-1. Table 2-11 shows the relationship between Figure 2-11 and Figure 2-12 and how they support the metrics in Table 1-1.

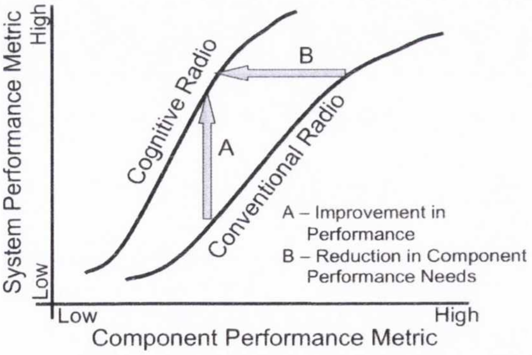

Figure 2-12 Assessment of Cognitive Radio Performance Benefits

Table 2-11 Computation of Cognitive Radio Benefit Metrics

	Mean Performance	**Reliability of Operation**
Improvement in Performance	Figure 2-12; Line A	Figure 2-11; *Point B/Point A*
Reduction in Required Hardware Capability	Figure 2-12; Line B	Figure 2-11; The difference in component performance between the cognitive radio lines

2.8 Extension of Methodology to Higher Levels of Network Operation

The framework described in this dissertation is directly extensible to the upper layers of a cognitive wireless device in a manner similar to the physical layers. A top-level view of a possible relationship of the physical and the network layer decision-making is proposed by Marshall [107] in DARPA's WNaN program, and is further discussed in Chapter 8. It demonstrates how limitations in the performance of the radio in specific environments and conditions are mitigated through adaptation in other areas of the performance region.

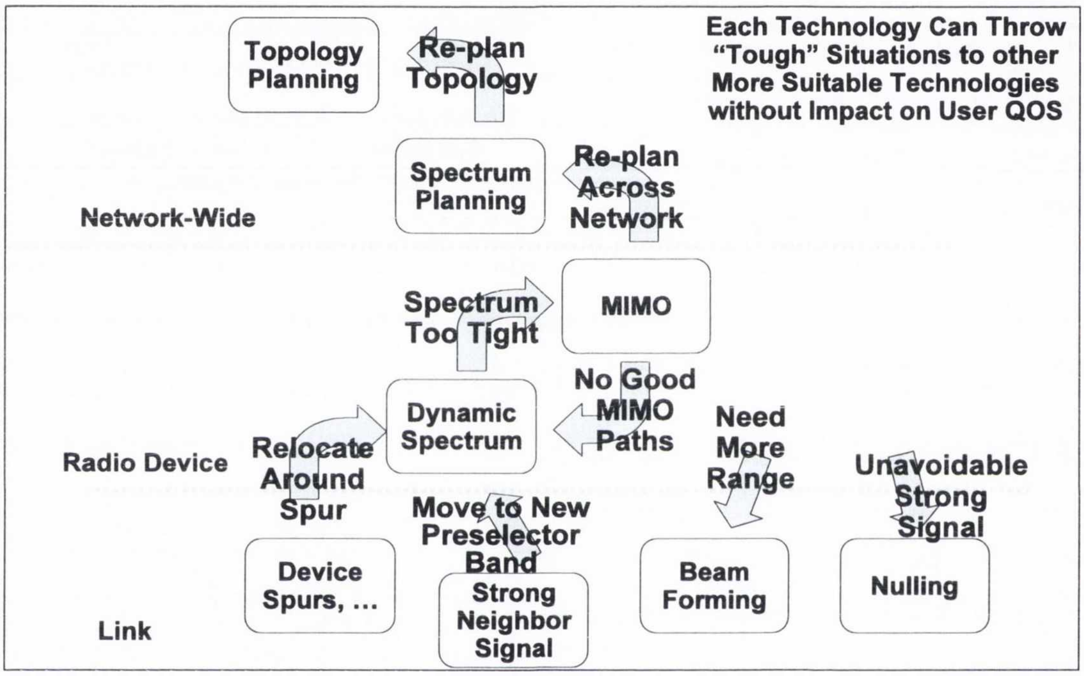

Figure 2-13 Relationship of Some Physical and Network Layer Decisions for a Cognitive Wireless Device

The lower layer of this depiction includes the adaptations whose performance and reliability effects are quantified by this dissertation. In the current PHY and MAC layer cognitive radio model, if spectrum is too scarce, there is little recourse. In the cognitive network model, however, this problem has potential redress, in that the network now has the opportunity to re-plan the topology to reduce spectrum requirements in the densely utilized region by modifying the traffic routing to regions that are more abundant in spectrum. For example, a network set up within a city might find that routing through the city was the shortest path, but created too much contention for spectrum resources. A less direct, but higher performing topology might move latency tolerant data out of the central core, around the perimeter, and back into the far side of the city! The need for this topology would only be apparent to a cognitive networking technology that had the degree of environmental awareness contemplated in this dissertation.

In this depiction, dynamic spectrum is the key linkage between the ability of the radio to adapt to address issues that arise within its physical environment, and the capability of the network to adapt to meet different and changing mission needs, within the fixed constraints of the environment and the radios basic performance capability. The basic emission tuple previously described is directly extensible to reflect these interactions with upper layer processes. A qualitative overview of these cognitive radio techniques is provided in Chapter 8.

This framework enables a network process to be concatenated onto the physical layer adaptation; reflecting the requirements of the network for various topologies, end-to-end latencies, packet error rates, and other metrics reflecting the upper layers of the network. Such an integrated process would provide arbitration between the link performance and network performance objectives. ***While much of the advance of the network field has benefited from clear abstractions of physical layer operations, this is not necessarily a virtue for highly resource-constrained systems***, such as: wireless networks; and precludes visibility into, and exploitation of options that could have significant benefit to the end users of such networks.

An example of such multi-layer integration is provided by Soon et al. [108] that reported the benefit of using MIMO to minimize interference between competing TCP transfers. TCP over wireless cannot distinguish expected traffic congestion with unexpected events, such as interference events and disrupted paths. This causes it to either declare the transfer a failure, or to reduce data rate, Even when the link and route return to operation, TCP may have reduced its offered load so that the link capacity is never fully exploited after the interruption event [109], [110].

In the chapters that follow, these principles will be expanded and quantified. In particular, it will be shown that these are strong cost advantages inherent in cognitive radio adaptations that can be an incentive for rapid adaptation of the technology; and, once a sufficient number of these devices are deployed, a second generation of DSA policies can be adopted that provide significant increases in device density through reliance on the inherent interference mitigation of cognitive radio.

Chapter 3 Cognitive Radio Environments

3.1 Introduction

This chapter provides an analysis and characterization of the spectrum environment in which a cognitive radio will operate. The different spectrum environments define the performance constraining regions that impact wireless systems operation, and must be addressed or mitigated by either intrinsic performance, or cognitive radio adaptations. Environments will be characterized both from empirical measurements, and by closed-form approximations that can be applied to generalize the results of the following chapters to any arbitrary environment. The use of closed-form expressions of spectrum environments will enable cognitive radio and DSA researchers to:

(1) Simulate a wider range of spectrum environments than can be sampled and analyzed;

(2) Perform analysis of radio performance, without researchers requiring large databases of environments; and

(3) Provide provable assertions about cognitive radio performance in a range of potential environments.

There have been a number of measures proposed to characterize the general occupancy of spectrum [111]. These are useful to assess the effectiveness of the spectrum management processes and policies, but are not adequate to assess the quantitative performance of devices that inhabit and exploit the characteristics of the spectrum.

The US National Science Foundation (NSF) sponsored a set of spectrum surveys in a number of city and rural environments. McHenry reported the data from these collections. The locations are shown in Table 3-1[1]. This data was selected as the basis for empirical examination of the spectrum due to: the comprehensiveness of the surveys, the technical and methodological consistency across collections at multiple sites, and their consistency with reports of other researchers. The technical details regarding the collection of each data set is provided in the referenced work. In general, the collections consist of a fixed antenna location, with collection durations of from 12 to 24 hours, and frequency scan rates from 30 to 120 seconds. A general description of the measurement campaign is provided by Bacchus et al. [112].

---

[1] Note that one site (NSF building roof in Arlington, VA) reported in the NSF study was not used in this analysis because the fixed antenna spectrum data above 1 GHz was reported to be unreliable, and therefore no comparable set of values could be obtained. Similar data previously collected from Chicago was used to provide a similar urban environment.

Table 3-1 Spectrum Collections

Sample	Location	Date(s)
Chicago	Illinois Institute of Technology, Chicago, IL [113]	November 16-18, 2005
Riverbend	Riverbend Park, Great Falls, Virginia [114]	April 7, 2004
Tysons	Tysons Square Center, Vienna, Virginia [115]	April 9, 2004
New York	Republican National Convention, New York City, New York (Day 1 and Day 2) [116]	August 30-September 2, 2004
NRAO	National Radio Astronomy Observatory (NRAO), Green Bank, West Virginia [117]	October 10 -11, 2004
Vienna	Shared Spectrum Building Roof, Vienna, Virginia [118]	Dec. 15-16, 2004

Some details on the environment of each sample are provided in Table 3-2.

Table 3-2 Spectrum Collections Environments

Sample	Details
Chicago	Building roof at periphery of high-density city environment, extremely elevated antennas with line of sight to most of the city.
Riverbend	A very small park along a parkway just above the Potomac River, just outside of Washington, DC. Although it is a park, it has close proximity to automobile traffic and associated emissions. A high amount of foliage attenuation in the direction of Washington, DC was present.
Tysons	A dense suburban shopping mall and office complex area outside of Washington, DC. The collection was at a shopping center along a highway, with considerable blockage of RF signals, but there were also local sources.
New York	Antennas located on the roof of Steven's Institute in New York, looking across the river to New York City. Line of sight directly to Manhattan Island, but at a considerable distance.
NRAO	Radio astronomy reservation protected from RF interference and emissions. Located in an extremely rural area of West Virginia. Quite distant from broadcast transmitters. Further information is available on their web site [119].
Vienna	Roof of tall office building in a suburban office park, within 10 miles of Washington, DC. Advantaged antenna location that overlooks Washington.

The sampled environments span the range of potential cognitive radio environments. Two of the locations are located in, or near the centers of major US cities (Chicago and New York). In contrast, one location, the National Radio Astronomy Observatory (NRAO) in Green Bank, West Virginia is one of the quietest RF locations on Earth, due to direct actions to eliminate potential RF sources from the environment. The Riverbend location had high foliage absorption, so the upper frequencies were highly attenuated, and therefore, quite vacant. The other sites were located within more urban environments, and have correspondingly higher signal densities and intensities. More detailed information on the measurement process at each location is available in the references. One site from the NSF Survey was not included due to lack of comparable spectrum coverage [120]. Note that none of the sites were in close proximity to high-powered emitters, so that the high energy (co-location effects) often encountered in practice by devices was not present in any of the samples. Therefore, they have a high probability of understating the peak energy to which a device immersed in a dense user environment might be subjected. In examining energy related performance (such as front-end overload), the measurements reported here should be considered to a lower bound on the likely energy intensity.

Although the measurement sets are technically quite consistent, there are some differences in the collection process. In some cases, no FM broadcast was recorded. FM broadcast is both a high power and relatively narrow signal, so its elimination from several of the measurement sets does potentially reduce the high-energy portion of the distribution, although it has minimal impact on overall spectrum occupancy and availability.

Section 3.2 develops the statistical characteristics of the spectrum energy within signaling bandwidths, with an emphasis on the low energy end of the distribution. These determine the spectral opportunities for cognitive radios to access spectrum. Experimental data collected in a range of spectrum environments is used to characterize the distribution of energy in each of the environments, and coefficients for a set of closed-form spectrum occupancy distribution functions are developed in order to generalize these results.

Section 3.3 describes the time-domain considerations of spectrum opportunities, including the possibility that the sensing operation will generate false detections that cause unnecessary spectrum relocations, and also investigates the statistical characteristics of DSA opportunities at a range of occupied spectrum thresholds.

Section 3.4 describes the occurrence and impact of high-energy conditions within the band-pass associated with the early receiver sections, typically a LNA and mixer. This discussion addresses the impact of total energy on the operation of linear RF devices, and the resultant effects on communications effectiveness. A generalized probability density function for high input energy is developed for each spectrum environment.

Section 3.5 generalizes the individual spectrum samples by creating a closed-form Cumulative Distribution Function (CDF) for a range of arbitrary spectrum environments, both for spectrum occupancy and high-energy regions. These functions enable arbitrary signal environments to be synthesized.

This discussion is intended to provide a set of empirically validated environments that will be encountered by a wireless device. Although specific values of environmental distributions are developed in this chapter, the analysis throughout the dissertation also provides closed-form representation of these distributions. Therefore, the results of this dissertation can be extended to a range of environmental characteristics, even if they are not reflected in the environment samples presented here.

Throughout the discussion of spectrum, two concepts of bandwidth will be described.

Signaling	Signaling bandwidth is the bandwidth that will be occupied by a transmitted or received signal, and is independent of carrier frequency. In this analysis, it can be considered to be the effective bandwidth of the last analog or filtering operation. It will be referred to as b_0 throughout this dissertation, and is an absolute measure (i.e., 25 kHz, 1, 2, or 10 MHz).
Pre-selector	Pre-selector bandwidth is the extent of the spectrum environment that is provided to the early active stages of the receiver, primarily the LNA. This bandwidth is considered to be a fixed ratio of the signaling frequency. It is referred to as Bandwidth Ratio (BW) throughout the dissertation, and is the effective octave range of the pre-selector. For example, a $BW = 0.2$ at 1 GHz would correspond to a 200 MHz pre-selector ranging from 900 MHz to 1.1 GHz, as will be further discussed in the text supporting Figure 3-27 in Section 3.4.3.

3.2 Spectrum Occupancy, Selection, and Access Characteristics

DSA has generally been postulated as the first fundamental benefit of cognitive radio, and has been the first that the wireless community has approached as a target for implementation. Early work by Mitola [5] pointed out that spectrum occupancy sensing was one of the obvious and potentially useful benefits of placing cognitive features within a radio. One important metric to determine is the probability that a given frequency, channel or range of frequencies will be available for use by a cognitive radio, consistent with the spectrum access policy controls that are imposed on the device's operation. For the purposes of this dissertation, this measure is referred to as *Spectrum Opportunity*, which is the ratio of frequencies usable by a cognitive radio at any given time, based on an energy threshold policy.

Spectrum Opportunity is one of the basic characteristics of the environment of a cognitive radio. As used here, it is the ratio of frequencies (at a given fixed bandwidth, b_0) that have power below the threshold compared to the total extent of frequencies. If b_o is equal to the smallest resolution bandwidth, then this function approaches the uncorrelated value of the distribution. As b_o increases, the probability of *Spectrum Opportunity* becomes influenced by the degree of correlation of the spectral openings, and is significantly impacted by the minimum required extent of openings. The complement of *Spectrum Opportunity is Spectrum Occupancy*, which is the probability that a frequency will be unavailable for assignment. Both metrics will be used to describe spectrum characteristics, as convenient.

A summary of the relative energy distributions of the spectrum samples is provided in Figure 3-1. The process used to analyze these will be discussed using the Chicago, IL sample set as an example, but the process is applicable to all sample sets; they differ only in parameter values, not in fundamental character.

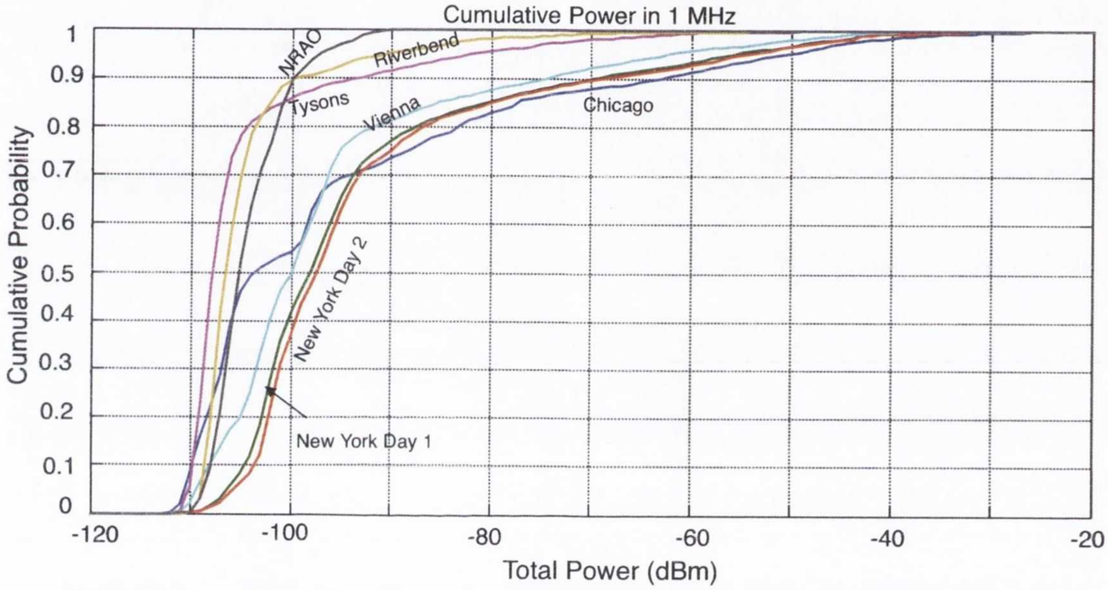

Figure 3-1 Summary of Spectrum Samples for 1 MHz Signaling Bandwidth (b_0)

This analysis characterizes individual spectrum time/frequency/energy profiles. Such a profile is shown in Figure 3-2; in this case Chicago from 30 MHz to 3 GHz approximately every 30 minutes for 24 hours.

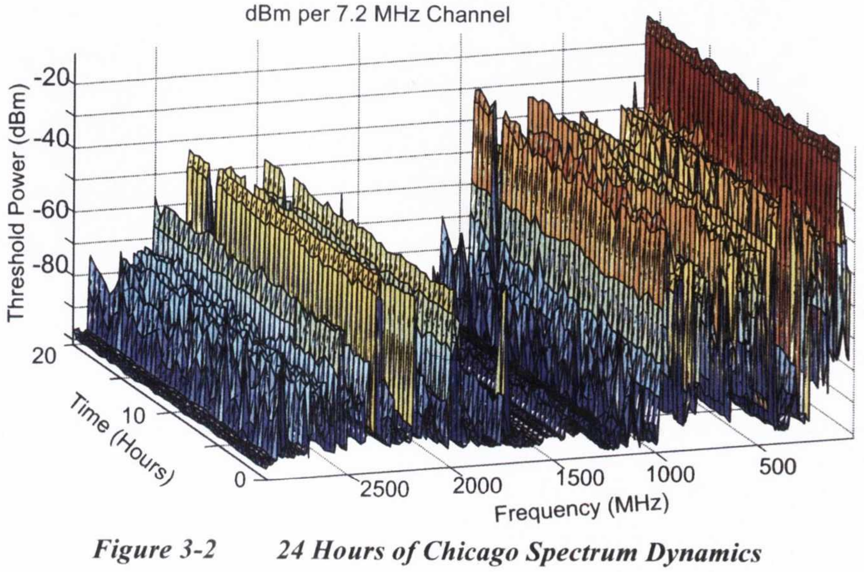

Figure 3-2 24 Hours of Chicago Spectrum Dynamics

Figure 3-3 illustrates the distribution of spectral energy as a function of bandwidth (b_0) and energy. This figure depicts the energy in contiguous bands from 25 kHz through 10 MHz and demonstrates the expected dependence of energy to the contiguous bandwidth.

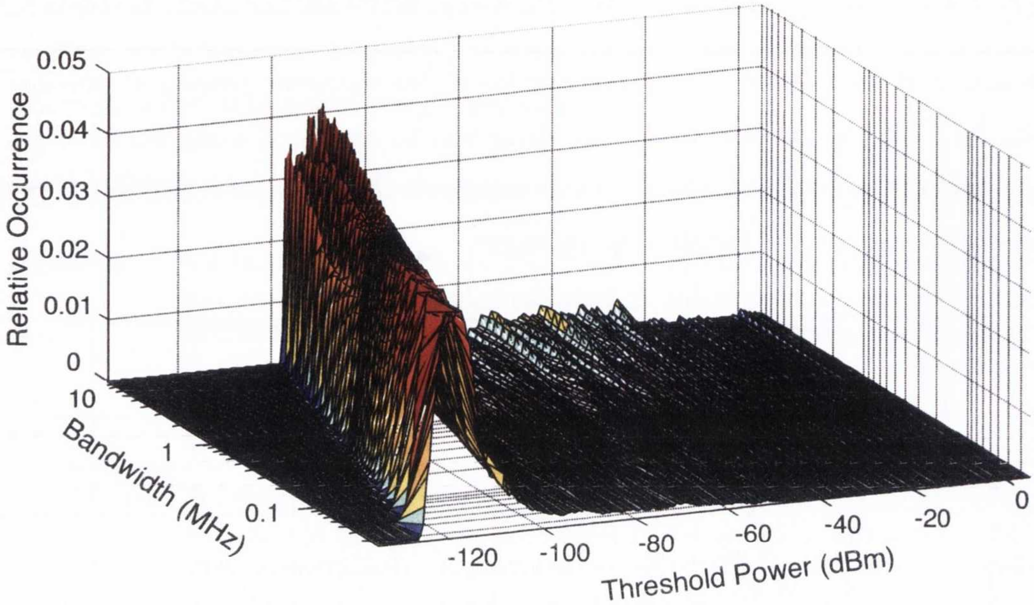

Figure 3-3 Spectral Distribution of Instantaneous Energy in the Chicago Spectrum Sample for a range of Signaling Bandwidths

These relationships of energy and bandwidth provide the mechanism to evaluate the operation of a cognitive radio within a given spectrum environment. The rest of this chapter will establish empirical measures of these relationships, and derive closed-form approximations for the spectrum environment of a cognitive radio.

The dependency of the *Spectrum Opportunity* measure on contiguous bandwidth (b_o) is important. Figure 3-4 illustrates the value of the Spectrum Opportunity for various ranges of b_o and threshold power in this same sample. Spectrum availability decreases rapidly with larger required contiguous bandwidth.

The Spectrum Opportunity ratio (for small values of occupied bandwidth) is shown in Figure 3-5 for a range of amplitude thresholds using mean values of the Chicago collection. These are essentially slices of Figure 3-4, and form the basis of a closed-form estimate of spectrum availability for a range of threshold and bandwidth parameters. As would be expected in a Gaussian environment, bandwidth values separated by a factor of two have noise threshold increases of approximately 3 dB.

The analysis to follow is performed as a function of signal and front-end bandwidth rather than the more generalized unit energy per unit bandwidth. A bandwidth independent metric, such as energy per Hz would be appropriate for independent and extremely narrowband signals, but would fail to reflect the correlated nature of the signals that will be encountered by a cognitive radio. If there is energy in one specific Hertz of the spectrum, it certainly increases the probability that there will be energy in the adjoining Hertz. This data shows that a threshold just at the noise floor results in the anticipated 100% occupancy (0%, or no opportunities), and that the probability of occupancy for high threshold values is low. A threshold 10 dB above the noise level shows 20% occupancy, for a 25 kHz signaling channel. The basic shape and character of the spectral density distribution is common to the entire set of spectrum samples investigated.

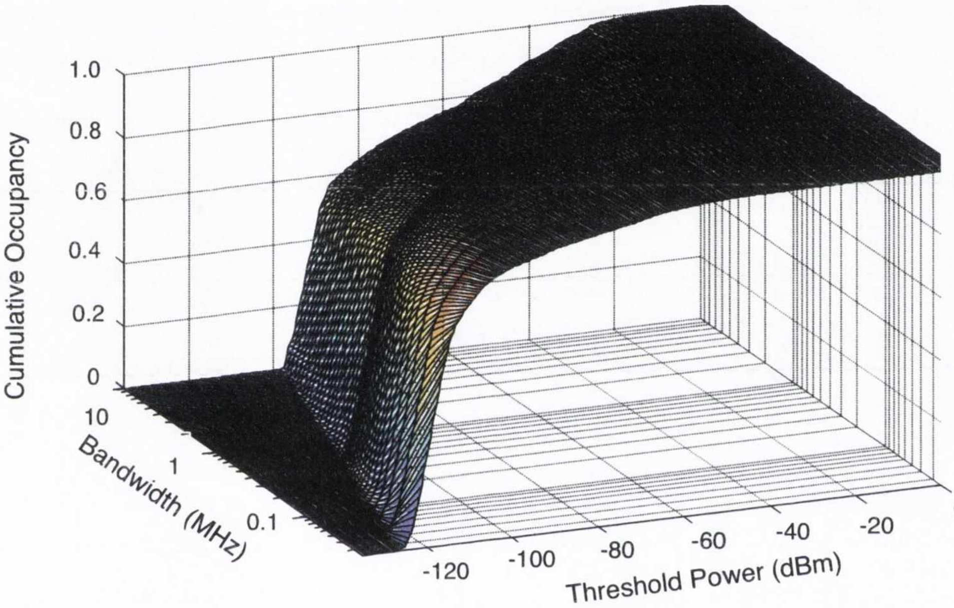

Figure 3-4 **Spectrum Opportunity for Chicago Collection Over a Range of Threshold and Signaling Bandwidth Values**

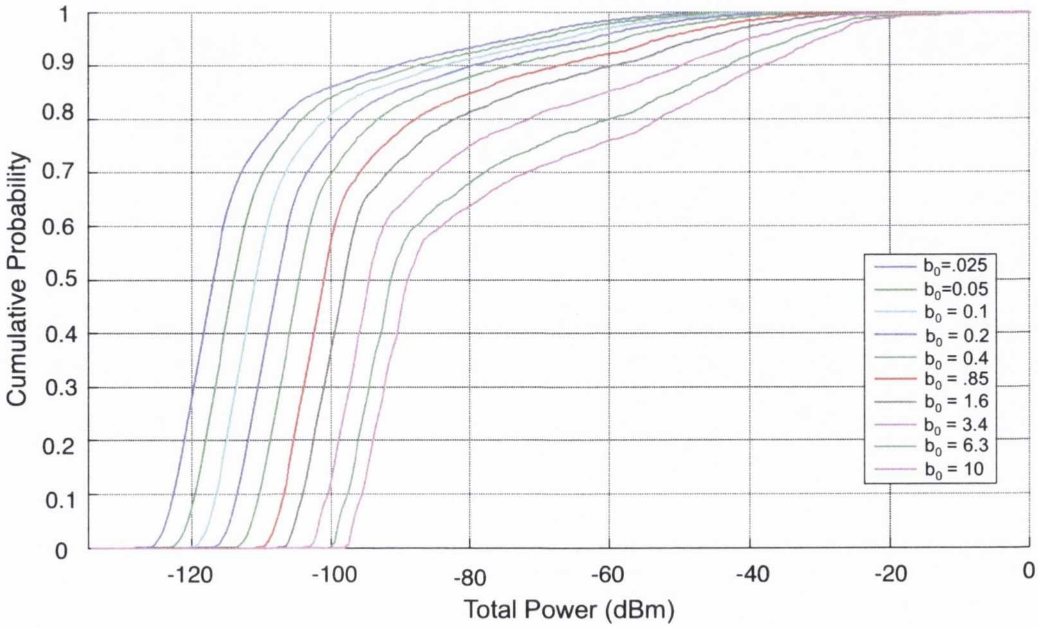

Figure 3-5 **Spectrum Opportunities Measurements for Chicago Spectrum as a Function of Bandwidth (b_0 in MHz)**

The distributions clearly demonstrate that the shape of spectrum occupancy is essentially indifferent to bandwidth (when sensed as an aggregate, integrated across the bandwidth). At low threshold values, the integrated bandwidth value (b_0) linearly displaces the intercept with the noise floor. The spectrum available at any given point is therefore a function of the level of the threshold above the noise floor and the noise bandwidth. As the threshold energy or bandwidth increases, the wideband or high-energy regimes exhibit the correlated effects discussed previously.

To create a representation of these energy probability distributions that is suitable for analytic treatment, a generalized form of the distribution is needed. The beta distribution was selected as this representation. The beta distribution is the conjugate prior of the binomial distribution. Its distribution is characterized by two parameters, α and β, which specify the degree of skewness and shape of the distribution around the mean.

The Beta distribution is the binomial distribution posterior distribution of parameter p after observing α-1 events with probability p and β-1 with probability 1-p. The beta distribution is not argued to be descriptive of any of the physical processes involved in spectrum density phenomena, but it has a number of unique advantages in this application. They include:

1. The beta distribution has zero probability of occurrence below a specific level, and also above a specific level. The tail of a Gaussian distribution would have non-zero probabilities beyond the measured data. Since communications links are designed around the probability of occurrence of relatively rare events, a distribution is needed that has zero probability of energy below the noise level, and correspondingly can reflect unity probability above the largest signal levels measured.
2. The Beta distribution has closed-form equations for the Probability Density Function (PDF) and Cumulative Distribution Function (CDF).
3. The distribution parameters are readily derivable from measured data characteristics (the mean and variance directly compute the α and β parameters).

The suitability of the beta distribution is also evident at the end of this section when the corresponding analytic and measured distributions are compared. Note that the Beta distribution is a reasonable engineering approximation, and that the variations in the Beta estimates are less than the differences between the experimental collections.

To determine the CDF of each spectrum sample distribution, each distribution is characterized as shown in Figure 3-6. The CDF is zero below and non-zero above the value of SO_{min}, which is specific to the spectrum distribution and bandwidth being considered, and is unity at all values above SO_{max}, which is also specific to the spectrum sample and bandwidth.

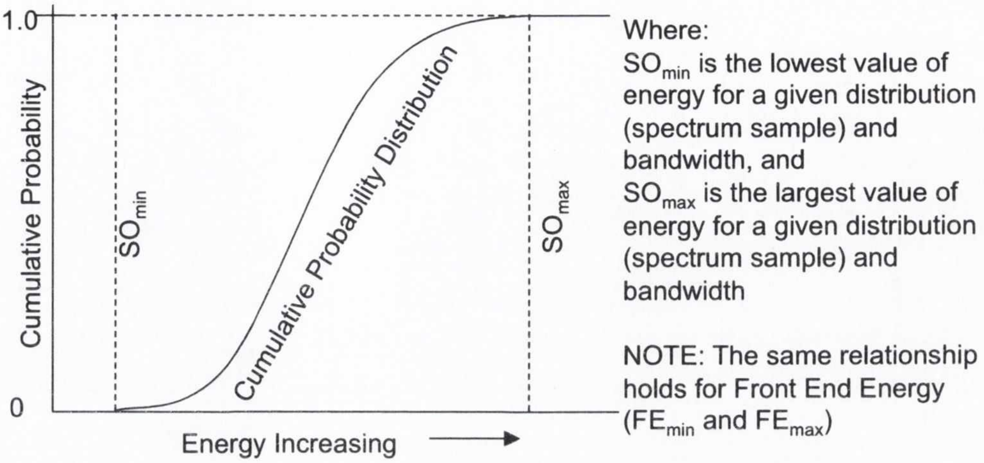

Figure 3-6 General Form of Spectrum Sample Characterization

The noise threshold levels ($SO_{\min}$) for the spectrum occupancy of each spectrum sample are shown in Figure 3-7. Considerable work has been reported on the determination and calibration of the noise floor of individual detectors, including both energy and feature detection [121], [122]. It is thus reasonable to assume that a detector can operate reliably in low energy conditions.

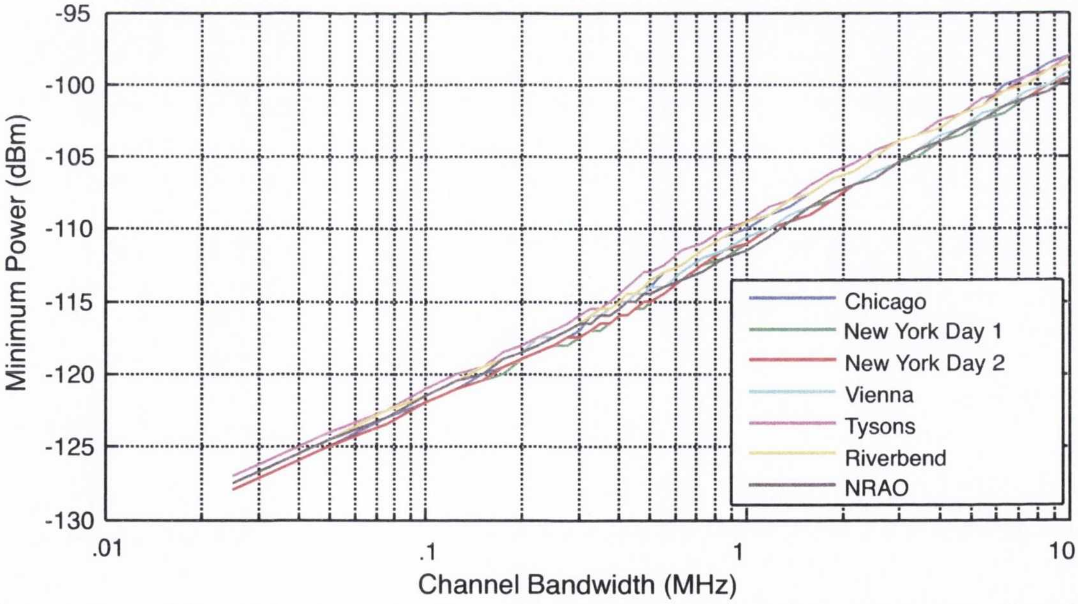

Figure 3-7 SO_{min} Noise Levels for Narrowband Signaling Bandwidths in Spectrum Samples

As would be expected, the lowest values of channel energy are decorrelated (noise dominated) and are essentially linear with the bandwidth, again reflecting that low energy values of each sample represent the noise bandwidth.

The corresponding upper energy limits ($SO_{\max}$) are shown in Figure 3-8.

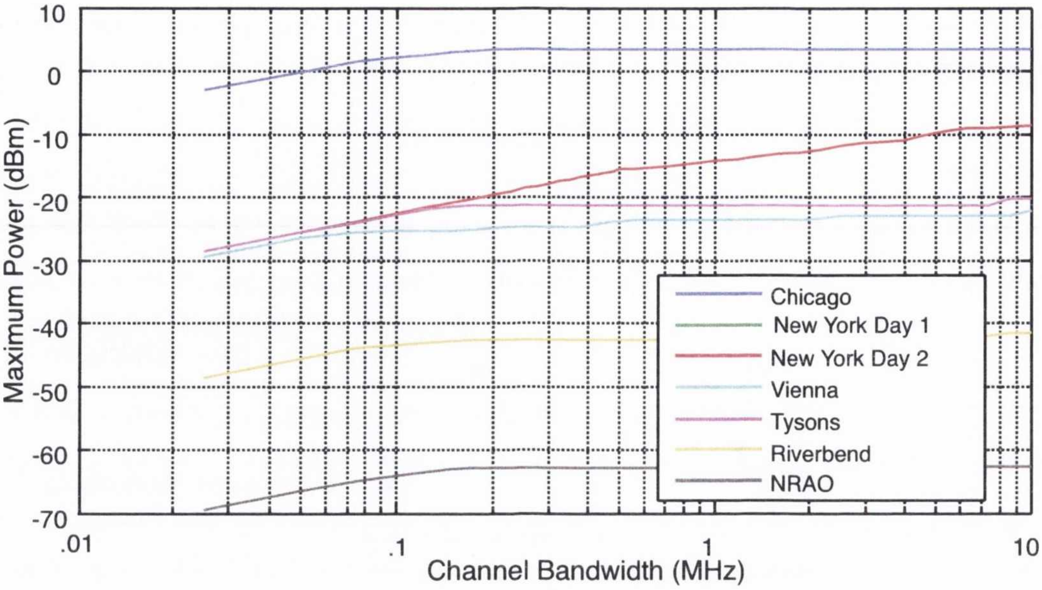

Figure 3-8 *SO_{max} Maximum Narrowband Energy Levels*

Several individual emitters in each sample dominate these values, and the results clearly reflect this correlation. Once the highest power emitter is entirely included in the pass-band, the increase in energy from integrating adjacent, but typically lower power channels is minimal. It is clear that high-density urban environments are quite distinct from rural ones, and that the differences between the urban measurements are not fundamental. The Chicago sample clearly reflects the fact that emitters are closer to the selected sensing site, and may be the most realistic representation of an urban location. The probability distribution between these two extreme points (the range between the SO_{min} and SO_{max} values) can be characterized by its statistical mean and variance. Figure 3-9 illustrates the mean value of this distribution.

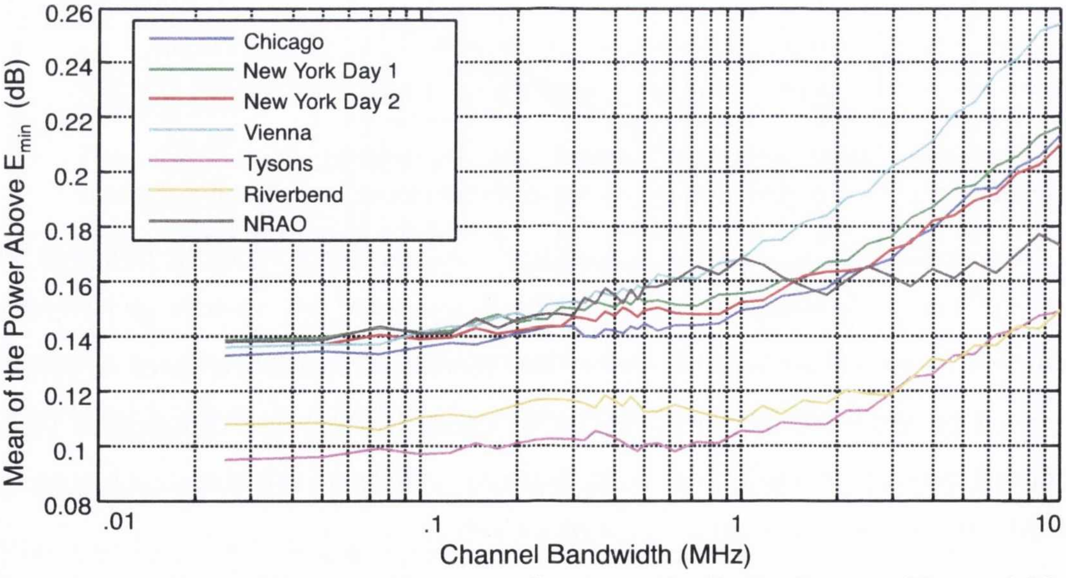

Figure 3-9 *Mean of Spectrum Occupancy Spectrum Distribution Between SO_{min} and SO_{max}*

Note that a small (less than 0.5) value indicates a skew towards the low energy signals, 0.5 is a symmetric (normal) distribution, and values above 0.5 indicate that the distribution is skewed towards higher energy levels. For the narrow band signaling case, all spectrum energy samples examined were strongly skewed to low energy populations. The mean values do not appear to cluster with any apparent density relationship. However, when plotted as absolute values *(mean (SO_{max}-SO_{min}));* they have a strong correlation with the signal density. This relationship is shown in Figure 3-10. The four dense environments all have absolute means above 14 dB, while the low density signals means are much closer to the minimum energy, typically less than 9 dB above the minimal energy levels.

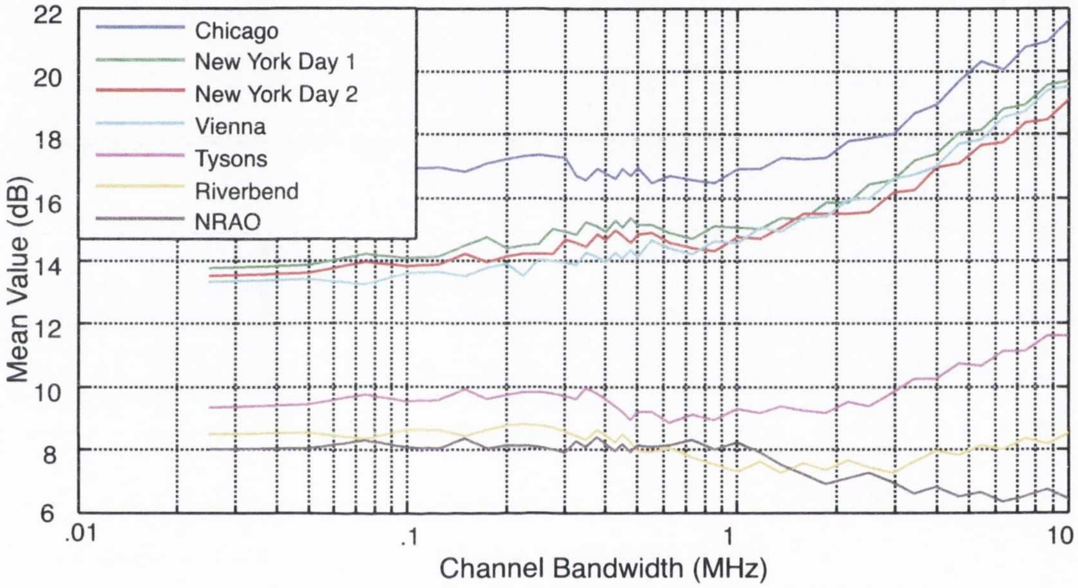

Figure 3-10 Absolute Mean of Signaling Bandwidth Distribution

The corresponding variance of the sample distributions is shown in Figure 3-11.

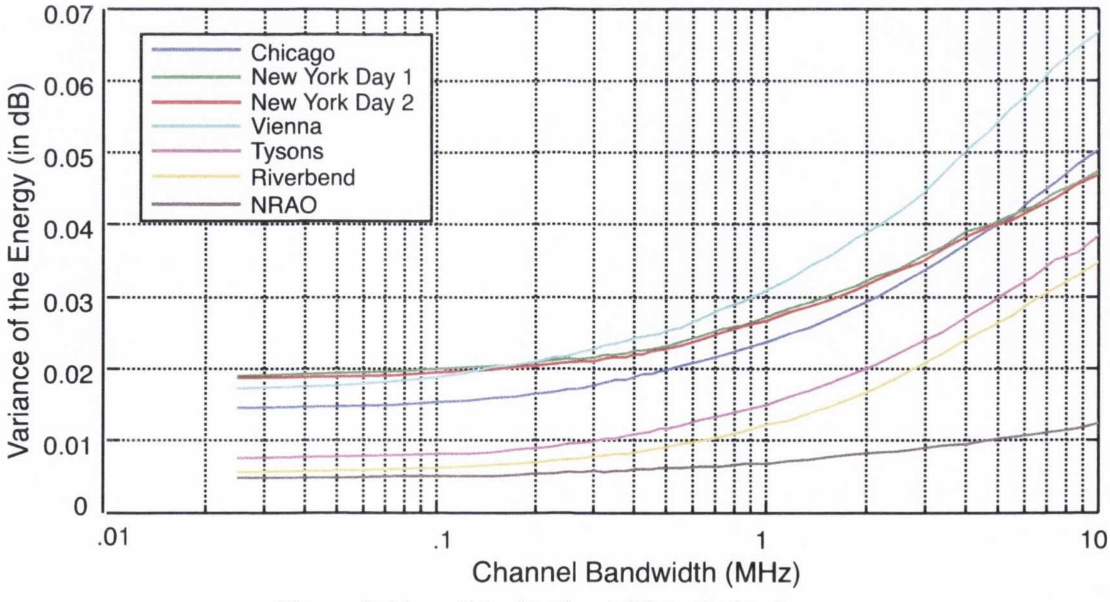

Figure 3-11 Distribution of Sample Variance

The ratio of the variance to the mean is provided in Figure 3-12.

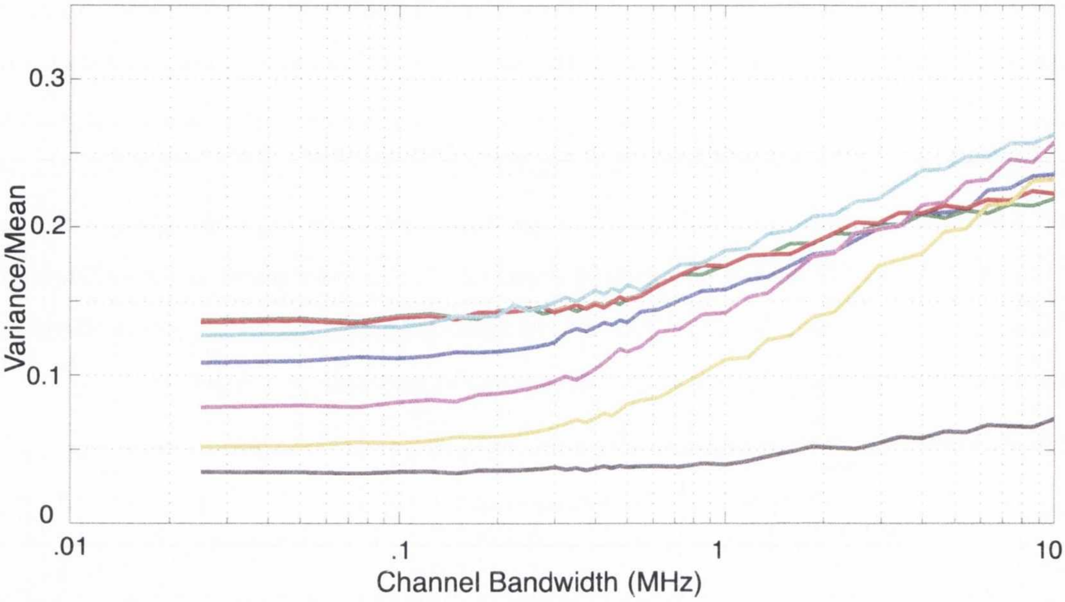

Figure 3-12 Distribution of Variance over the Mean

Again, the pattern of increased energy extremes with increased spectrum density continues. The three lowest density environments have a variance that is typically one half that of the higher density environments, through this same range.

The beta distribution α and β parameters are derivable directly from the mean and variance, and are shown in Figure 3-13 and Figure 3-14 respectively [123].

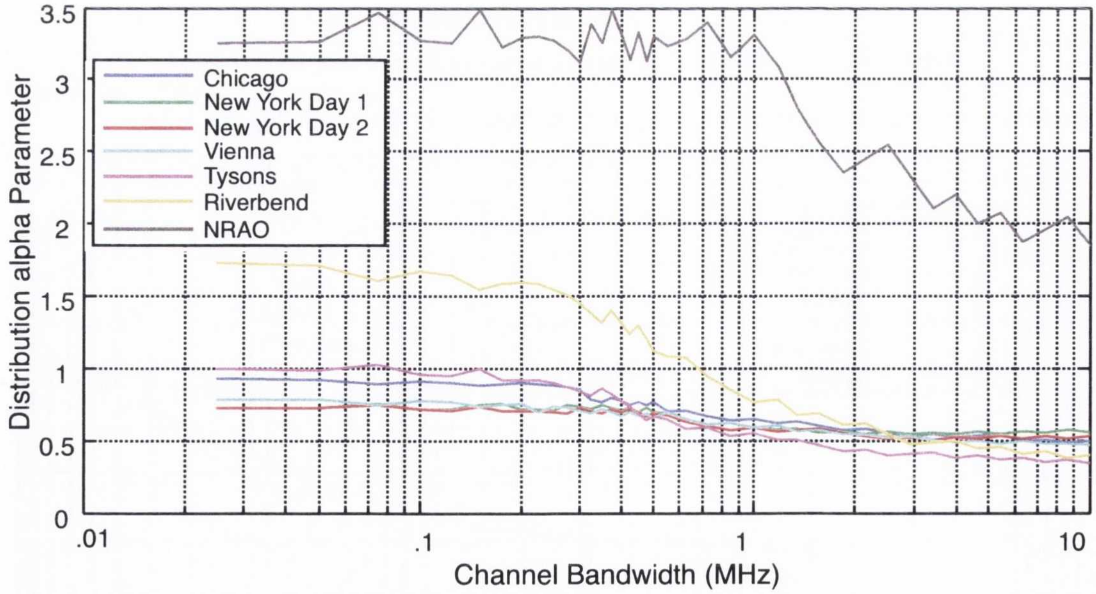

Figure 3-13 Spectrum Occupancy Beta Distribution alpha Value for Spectrum Samples

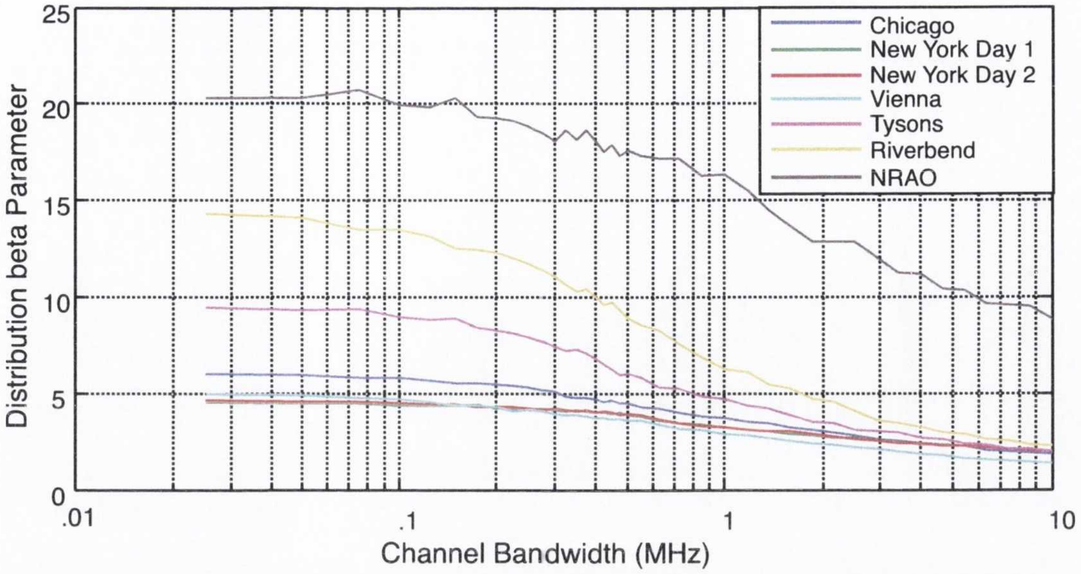

Figure 3-14 Spectrum Occupancy Beta Distribution beta Value for Spectrum Samples

Again, the alpha (α) parameter values of the NRAO distribution reflect the correlation and scarcity of the emitters in the environment. For the stressing, dense environments of interest, these measures are relatively constant. The beta (β) parameter is clearly monotonic to the overall density of the environment. As the environment becomes denser, the "law of large numbers" dominates: the distribution is less Poisson, and the beta reduces. The same effect occurs as the bandwidth is increased.

The values of the minimum and maximum energy, and α and β distribution parameters are usable directly from these graphs to characterize environments and understand performance sensitivities, but it is also useful to develop a closed-form expression to characterize arbitrary spectrum environments and signaling bandwidths. Each of the parameters is modeled as a closed-form equation of the channel bandwidth (b_0). These equations are shown below:

$$SO_{\min} = k_{SOmin2} \log_{10}(b_0)^2 + k_{SOmin1} \log_{10}(b_0) + k_{SOmin0} \text{ (in dBm)} \quad (3\text{-}1)$$

$$SO_{\max} = k_{SOmax2} \log_{10}(b_0)^2 + k_{SOmax1} \log_{10}(b_0) + k_{SOmax0} \text{ (in dBm)} \quad (3\text{-}2)$$

$$SO_{\alpha} = k_{SO\alpha2} \log_{10}(b_0)^2 + k_{SO\alpha1} \log_{10}(b_0) + k_{SO\alpha0} \quad (3\text{-}3)$$

$$SO_{\beta} = k_{SO\beta2} \log_{10}(b_0)^2 + k_{SO\beta1} \log_{10}(b_0) + k_{SO\beta0} \quad (3\text{-}4)$$

The *SO* terms are the characteristic parameters of the Beta signaling bandwidth energy distribution. The empirically derived polynomial coefficients ($k_{SO...}$) for the estimators of the sample sets are derived by least squares regression, and are shown in Table 3-3 and Table 3-4.

Table 3-3 Empirically Derived Spectrum Occupancy Distribution Constants

	CDF α			CDF β		
Sample	$k_{SO\alpha 0}$	$k_{SO\alpha 1}$	$k_{SO\alpha 0}$	$K_{SO\beta 2}$	$k_{SO\beta 1}$	$k_{SO\beta 0}$
Chicago	0.683	-0.227	-0.0147	3.87	-2.1	-0.287
NY Day 1	0.641	-0.117	-0.0184	3.4	-1.35	-0.314
NY Day 2	0.622	-0.129	-0.0216	3.41	-1.39	-0.314
Vienna	0.62	-0.163	-0.0237	3.05	-1.74	-0.203
Tysons	0.604	-0.342	0.00323	5.11	-3.74	-0.126
Riverbend	0.961	-0.716	-0.0508	7.27	-6.05	-0.305

Table 3-4 Empirically Derived Spectrum Occupancy Energy Constants

	Signal Ceiling (SOmax)			Signal Ceiling (SOmax)		
Sample	k_{Somax2}	k_{Somax2}	k_{Somax2}	k_{Somax2}	$k_{So\alpha 1}$	$k_{So\alpha 0}$
Chicago	-110	12.2	0.481	3.83	0.495	-1.74
NY Day 1	-111	11.5	0.629	-14	6.75	-1.54
NY Day 2	-111	11.6	0.652	-14	6.75	-1.54
Vienna	-111	11.3	0.507	-23.6	1.5	-0.864
Tysons	-110	11.6	0.26	-20.7	0.801	-1.88
Riverbend	-110	11.8	0.391	-42	0.831	-1.23
NRAO	-111	11.2	0.598	-62.3	0.631	-1.72

These polynomial coefficients make it possible to interpolate or extrapolate the spectrum measurements to different signaling bandwidths or environmental energy levels. To obtain an estimate of how these polynomial expressions reflect the empirical distributions, the empirical and closed-form estimates of minimum and maximum energy are illustrated in Figure 3-15, which illustrates the fit of the closed-form estimates of the minimum and maximum energy parameters. The curves are the closed-form estimates obtained from the polynomial expression from Equations (3-1) – (3-4) applied to the Chicago measurement set.

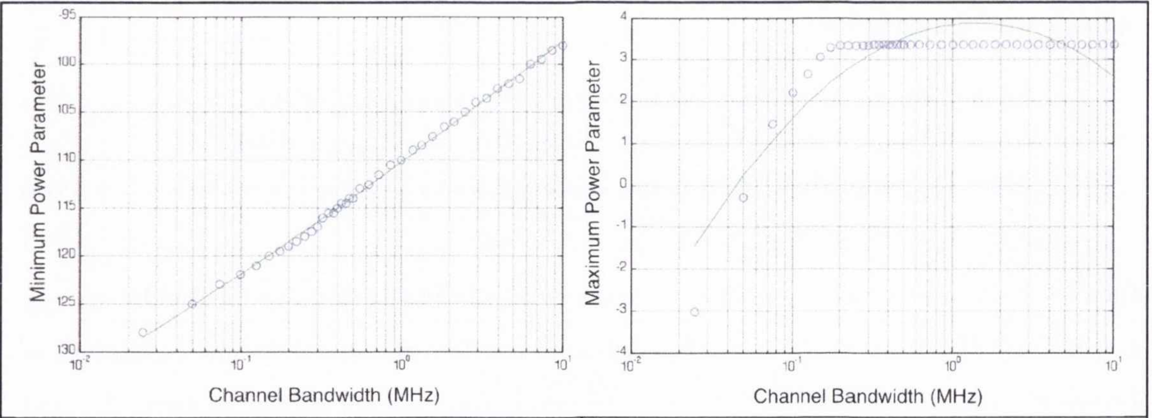

SO_{min} Parameter **SO_{max} Parameter**

Figure 3-15 Empirical and Modeled Distribution for Chicago minimum and maximum Energy

Similarly, the fit for the alpha and beta parameters are provided in Figure 3-16.

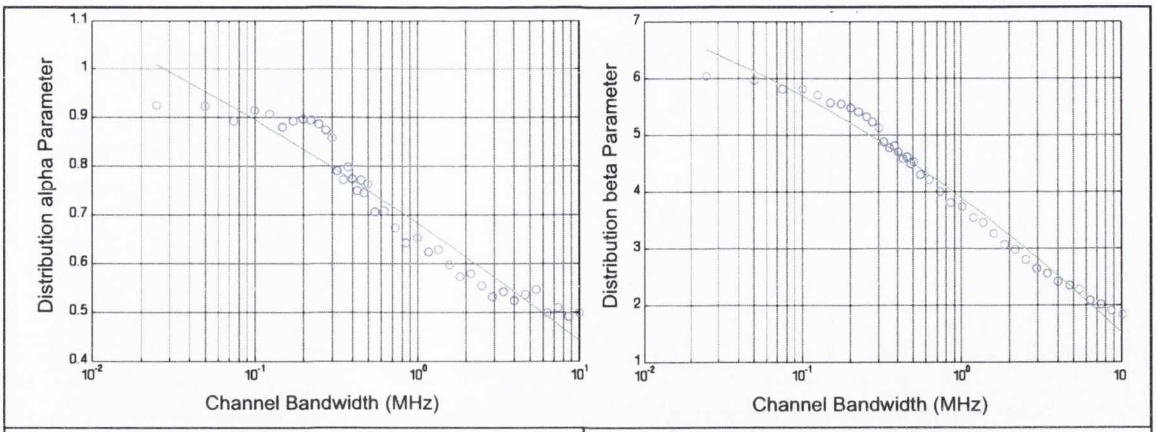

Beta Distribution α Parameter	**Beta Distribution β Parameter**

Figure 3-16 Empirical and Closed-form Distribution Parameters for Chicago Spectrum Occupancy Function

Note that while the polynomial estimates are not "perfect", the variations are small compared to the differences in the spectrum collections.

The cumulative distribution function of a Beta distribution is computed by:

$$CDF(x;\alpha,\beta) = I_x(\alpha,\beta) = \frac{\beta_x(\alpha,\beta)}{\beta(\alpha,\beta)} \tag{3-5}$$

Where:

$B_x(\alpha,\beta)$ is the incomplete Beta function
$B(\alpha,\beta)$ is the Beta function
$I_x(\alpha,\beta)$ is referred to as the regularized incomplete Beta function.

The Beta function in the denominator normalizes the incomplete Beta function to an appropriate probability range from zero to unity.

The CDF of the spectral energy is therefore given by:

$$SpectrumOpportunities(a_t,b_0) = \frac{\beta_x(SO\alpha,SO\beta)}{\beta(SO\alpha,SO\beta)} \tag{3-6}$$

Where:

a_t is the amplitude threshold

$$x = \frac{a_t - SO_{\min}(b_0)}{SO_{\max}(b_0) - SO_{\min}(b_0)}$$

To obtain a sense of the accuracy of these estimators, Figure 3-17 provides the empirical and estimated occupancy estimates for the Chicago spectral energy distributions.

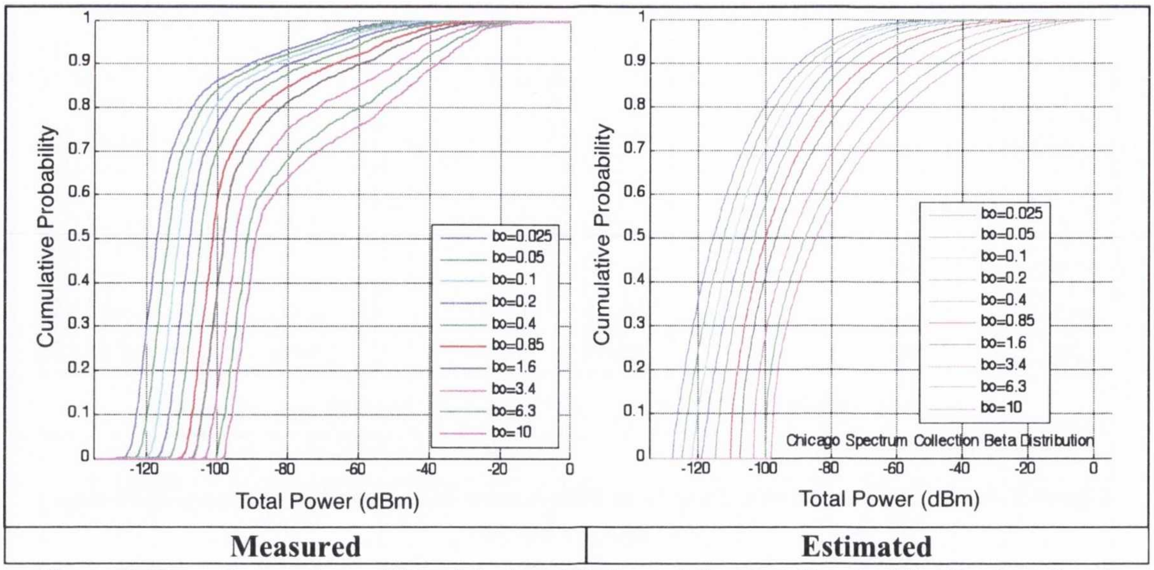

Measured	**Estimated**

Figure 3-17 Comparison of Empirical and Estimated Spectral Density Distributions for the Chicago Spectrum Sample

The benefits of this expression are twofold. First, it provides a convenient and tractable composite of over 7×10^8 individual data points. More important, it will provide a closed-form framework in which assertions about cognitive radio performance can be more universally proven for a class of environments, rather than demonstrated in a single environment. A summary of the modeled spectrum distribution for a narrow signaling bandwidth is shown in Figure 3-18.

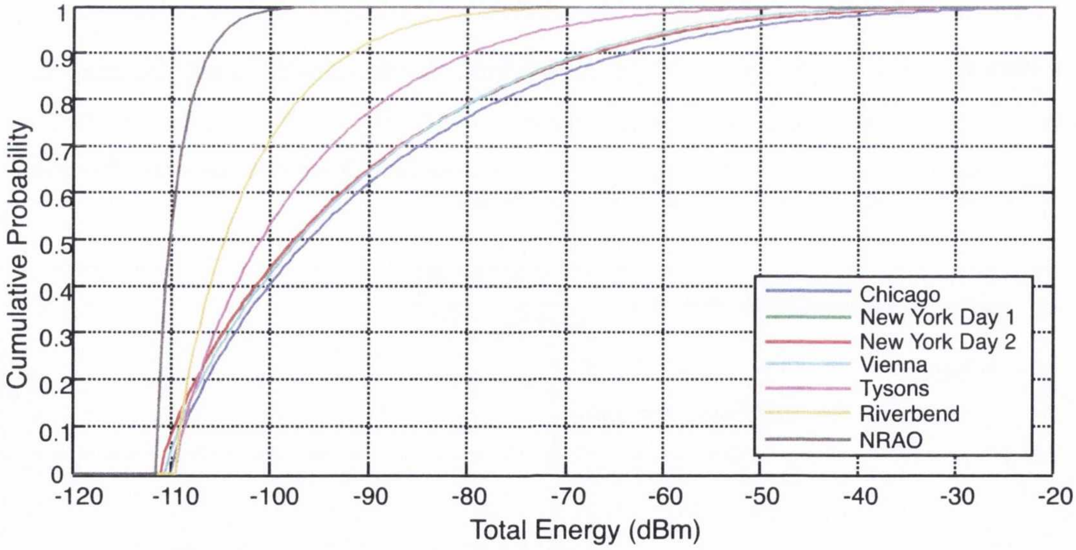

Figure 3-18 Parametric Distributions of Spectrum Samples at 1 MHz Bandwidth

3.3 Time Varying Characteristics of Spectrum Opportunities

The second aspect to be considered is the time duration of individual opportunities. Using the same data sets, the sensitivity of spectrum openings to the time duration of the opening is shown in Figure 3-19.

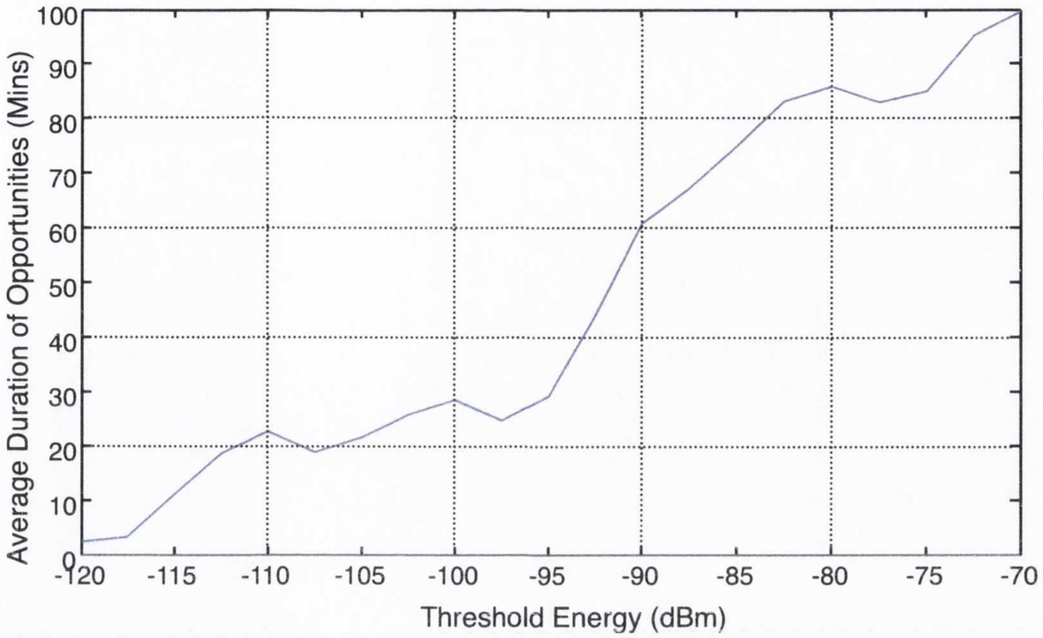

Figure 3-19 Stability of Spectrum Opportunities for a Variety of Duration Values in the Multiple Spectrum Collections and Predicted False Alarm Rate

Openings are much less stable at low values of occupancy threshold, while at higher threshold levels, there is a reduced likelihood of an opening being terminated. For example, at a threshold range of 5 dBm above the noise, the sensor reports above threshold power on an average every 10 sensor scans (minutes at the schedule of the sensor data provided), which corresponds approximately to the false alarm rate predicted by the discussion that will occur later. Until the threshold is raised to around 10 dB above noise, the duration of opportunities is short. This would not appear to be characteristics of the spectrum environment, but is a floor on the estimation provided by the intrinsic false alarm rate of the sensing process. The false alarm rate dominates the sensing at low thresholds, and the true duration of spectral openings is effectively masked by this unavoidable constraint.

The previous discussion assumed that the detection probability had perfect fidelity to the actual RF environment; any detected signal was due to the presence of a true interferer, and that lack of a signal indicated the lack of an interferer. The validity of this assumption falls into two measures, shown in Table 3-5, and reflect the Probability of False Alarm[1] (P_{FA}) and Probability of Detect (P_D) respectively.

---

[1] The P_{FA} and P_D measures are from the perspective that the "target" is a signal of sufficient energy to meet the spectrum policy threshold, and is detected (P_D) or that is reported when it is not present (P_{FA}). Probability of False Alarm (P_{FA}) is not a regulatory concern, and only has performance impacts, while P_D failures could result in causing interference to the source.

Table 3-5 Spectrum Sensing Detection Performance Categories

	Term	Cause	Effect
P_{FA}	Probability of False Alarm	Total random noise levels exceed the energy threshold by which the channel is declared to be occupied, and the channel is falsely declared occupied.	Forces the cognitive radio to change frequencies unnecessarily, or shows a properly available frequency as occupied, when in fact it is not.
P_D	Probability of Detect	1-P_D is the probability that a signal that is present is not detected, due to the normal variations of signal + noise energy detected.	Causes the cognitive radio to declare a frequency usable when it is not, potentially causing interference events for both the cognitive and non-cognitive radio.

The P_D is mostly of concern to regulators, as the negative effects of a detection failure will be felt by the protected user, not the cognitive radio. However, the P_{FA} is of concern to the cognitive radio, since it forces the radio to invest energy and channel time in relocating networks unnecessarily, and causes a loss of spectrum opportunities due to incorrect elimination of candidate frequencies. Energy detection of signals over short scan times is similar to amplitude detection of radar pulses and follows the classic radar Receiver Operating Curve (ROC) equation by Skolnik [124]. The P_{FA} is shown in Equation (3-7) and illustrated in Figure 3-20.

$$P_{FA} = \exp\left(\frac{a_t}{2}\right) \tag{3-7}$$

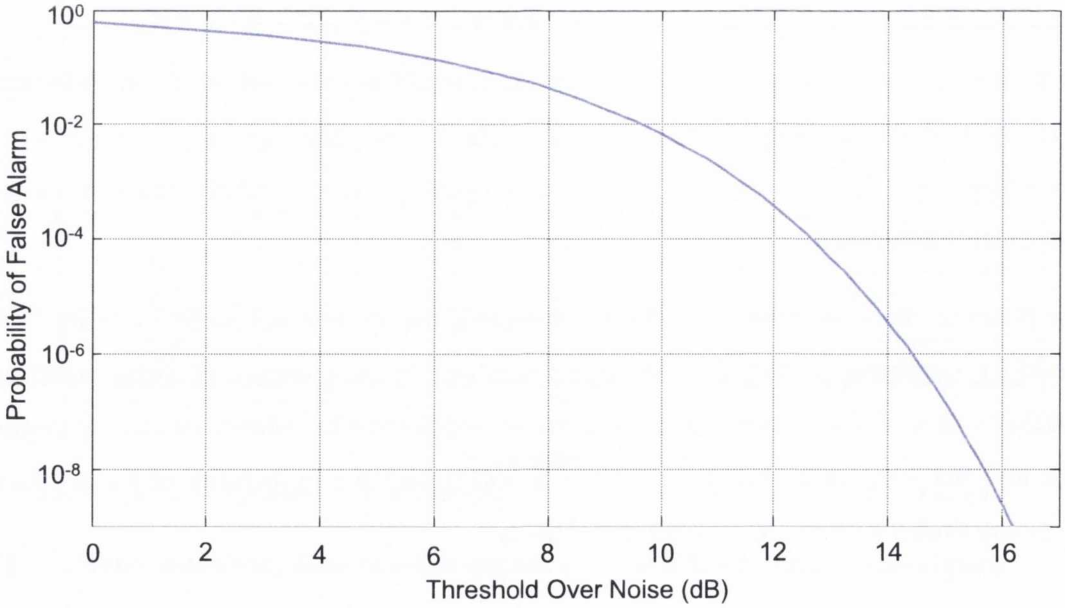

Figure 3-20 Illustrative Values of P_{FA} for a Range of Detection Thresholds

These effects are not significant for a detection threshold of at least 12 dB above the noise floor, which depends on the degree of non-interference guarantee that must be provided to incumbent users. Another way of looking at this parameter is to compare the spectrum available in the spectrum surveys to the false alarm rate of the corresponding thresholds. This measures the proportional loss of spectrum access for an incremental reduction in P_{FA} (corresponding to increasing the occupied threshold). An example of this relationship is shown in Figure 3-21. The similar shape of the P_{FA} curves is due to the relatively constant slope of the spectral noise distribution above the noise threshold.

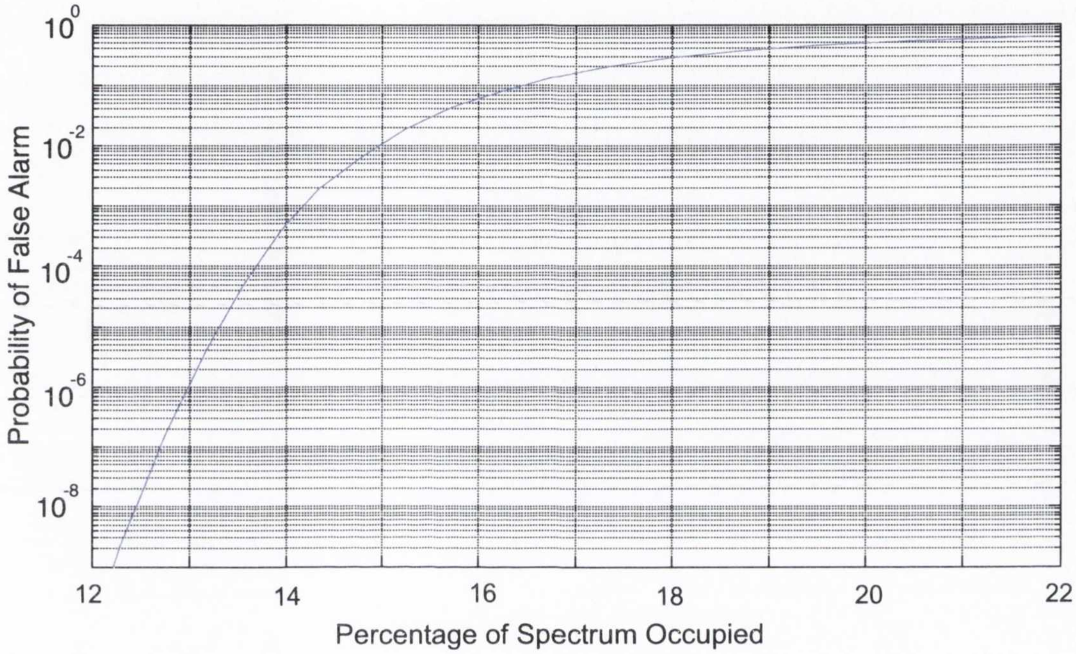

Figure 3-21 Spectrum Occupancy as a Function of P_{FA} for Chicago Spectrum

There is a very significant dependence of the aggregate network False Alarm rate on the threshold selection criteria, the number of scans, the resolution bandwidth, and with the size of the network (over which all nodes must agree on the occupancy of the selected frequency). In many cases, these effects scale with the number of nodes.

From this graph, it is evident that the false alarm rate can be reduced by three orders of magnitude at the cost of ignoring around one percent of the spectrum[1]. This implies that (for these environments) a cognitive radio can achieve essentially arbitrary false alarm rate without significant spectrum access loss. It is therefore a reasonable assumption that the cognitive radio could utilize two thresholds; one for deciding to enter the frequency, and one (established by interference criteria) that dictates exiting from the frequency. Even a few dB of difference between these thresholds can provide essentially arbitrary immunity

---

[1] For example the P_{FA} at a threshold that results in 13% of the spectrum occupied is 10^{-6}, while at a threshold resulting in 14% occupied spectrum, the P_{FA} level is approximately 10^{-3}.

to False Alarms, at the cost of a few percent of spectrum access[1]. Additional confidence that the false alarm problem is not fundamental is provided by Zhao and Swami, who prove that a decision-theoretic approach to spectrum sensing and opportunity identification enables these two aspects of a cognitive radio to be treated as separable problems, and therefore they can be optimized individually [125].

The gap allowed between these thresholds will be referred to as entry/exit hysteresis. A comparison of the spectrum cost of providing various levels of (Gaussian environment) false alarm (via entry/exit hysteresis) mitigation is shown in Figure 3-22. An example reading of this chart is that if an additional 8% of the available spectrum is made invalid for entry (non-enterable) because the energy was above the entry criteria, but still valid for continuing operation (because the energy was still below the regulatory threshold that would force exit), then the P_{FA} will be approximately 10^{-4}. By this mechanism, small perturbations in the signal detected at the cognitive radio will not cause a rapid change in status of a given channel between valid and invalid states.

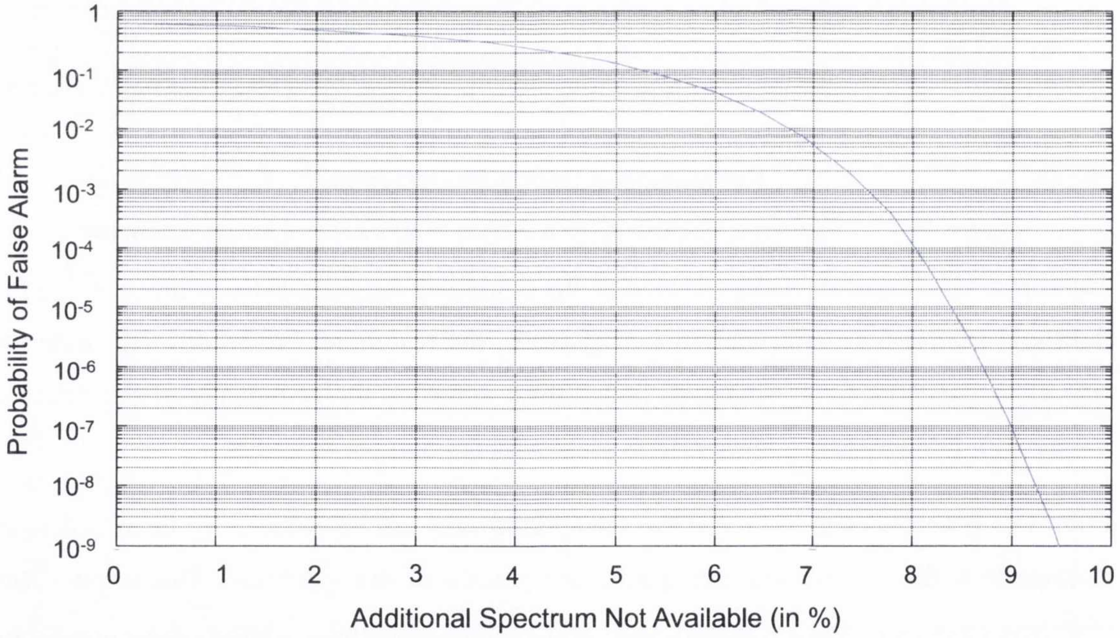

Figure 3-22 Lost Spectrum Access as a Function of Achieved P_{FA}

In this figure, the "Additional Spectrum Not Available" axis refers to spectrum that meets the regulatory criteria for valid use, but has noise in excess of the entry criteria (noise level is within the hysteresis region), and therefore is not available for initial entry, but operation could be permitted to continue once entered. There is no reason that these

---

[1] In practice, an additional consideration would be the range of fading events that the link to the interferer may be subject to. These correlated effects may be more significant and less predictable than the Gaussian ones in the figures.

criteria must be fixed or predefined; it would be most advantageously applied if it adjusted to reflect the relative scarcity of spectrum.In this dissertation, discussion of dynamic spectrum generally focuses on the ability to sustain a medium to long-term occupancy of channels. Shorter term sharing, such as an interleaved MAC Layer is also a spectrum sharing opportunity, and is discussed by Geirhofer et al. [126].

Consider two simplistic environments; one with stable broadcast uses which change only very infrequency, and one with fast frequency hoppers that constantly cycle through the available frequencies. In both cases, they may have identical *Spectrum Occupancy* functions, but the ability to exploit the latter is clearly much more constrained. In current practice, the spectrum usage is dominated by stable signals, resulting in stable opportunities. The deployment of significant numbers of cognitive radios will change that environment, and is an important issue for future dense deployments of cognitive radios.

3.4 Front-end Energy Distributions and the Effects of Front-end Non-Linearity

3.4.1 The General Problem of Front-end Linearity

The previous sections discussed spectrum occupancy from the premise that each frequency operated independently of every other frequency, and that the process of examining spectrum was separable, in that each individual frequency was an independent decision. This section addresses the practical effects and constraints introduced by realistic device operating characteristics and performance.

One of the most important effects is the impact of total input energy on the performance of the receiver front-end. Advocates of DSA have often assumed that any available "white space" would be usable for communications. In fact, white space next to a strong emitter may be un-exploitable, regardless of policy. Section 3 examined the energy environment for small signaling bandwidths; this section examines the typically larger front-end pre-selector bandwidths that drive the intermodulation response of the radio front-end circuitry. Another important distinction is that pre-selector operation is generally proportional to the carrier frequency, whereas Section 3.2 considered the signaling bandwidth as a constant bandwidth[1] regardless of center frequency.

In fact, the drive for addressing these concerns is not specific to DSA systems. In the US, a lengthy and controversial regulatory proceeding was initiated to resolve interference

---

[1] Such as in a typical fixed bandwidth Intermediate Frequency (IF) filter.

between the base stations of a cellular service provider, NEXTEL, and numerous local public safety systems. In the end, the US regulator (the FCC) elected to relocate the cellular systems through a mix of spectrum offerings, and cellular provider contributions to Public Safety frequency relocation [127]. These systems did not overlap in frequency usage, so this was not a frequency management failure, as typically defined. But the placement of high power cellular base stations did have a very significant impact on the performance of the public safety radio systems due to the very high energy level in adjacent frequency bands. Similar anecdotal experience is often referred to as "Co-site interference", as it is often the product of placement of a receiver in close proximity to a relatively strong emitter within the frequency response range of the receiver's front-end.

One of the contentions of this dissertation is that front-end overload, due to adjacent channel energy,[1] is a common and generally under-recognized experience. With denser spectrum assignments, and, as technologies such as DSA become more prevalent, this phenomenon will become a fundamental constraint on wireless networking. Its solution is therefore imperative if wireless density is to be increased, as proposed by DSA advocates, with reasonable constraints on receiver front-end linearity, cost, and energy consumption. As RF devices become integrated, the inherent limitations of CMOS processes will result in the future highly integrated RFICs being even more sensitive to the effects of constrained front-end linearity [128]. A design analysis of an SDR based on the New York City samples is provided by Hickling et al. [129].

A representative quantitative example of this situation is shown in Figure 3-23. Figure 3-23(a) depicts the input spectral distribution of eight (unequally spaced) signals as a typical input to an LNA. In this case, the total energy is just below the IIP3 of the amplifier. The input signals have separations ranging from 40 to 100 MHz. Figure 3-23(b) illustrates the spectral distribution of the same signal set after amplification by the receiver front-end LNA. The LNA processing these input signals generates intermodulation products or artifacts essentially every 10 MHz throughout the octave range, and beyond. A denser and more complex signal mix would create correspondingly denser, and even more complex sets of intermodulation products.

---

[1] As compared to the intended signal, this is in-band and can be addressed through gain control or attenuation without significant impact on the reception quality.

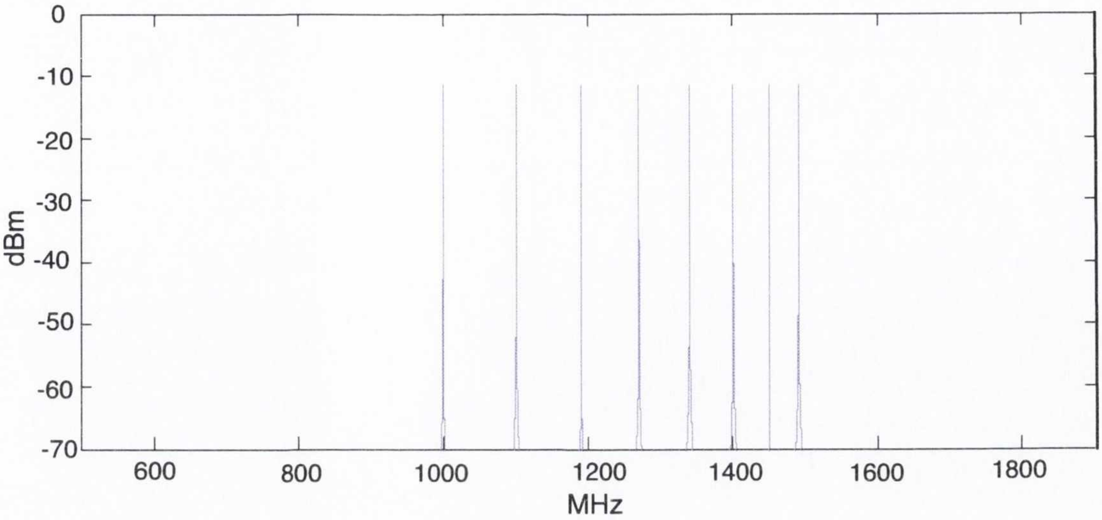

(a) Input Signal Spectral Distribution

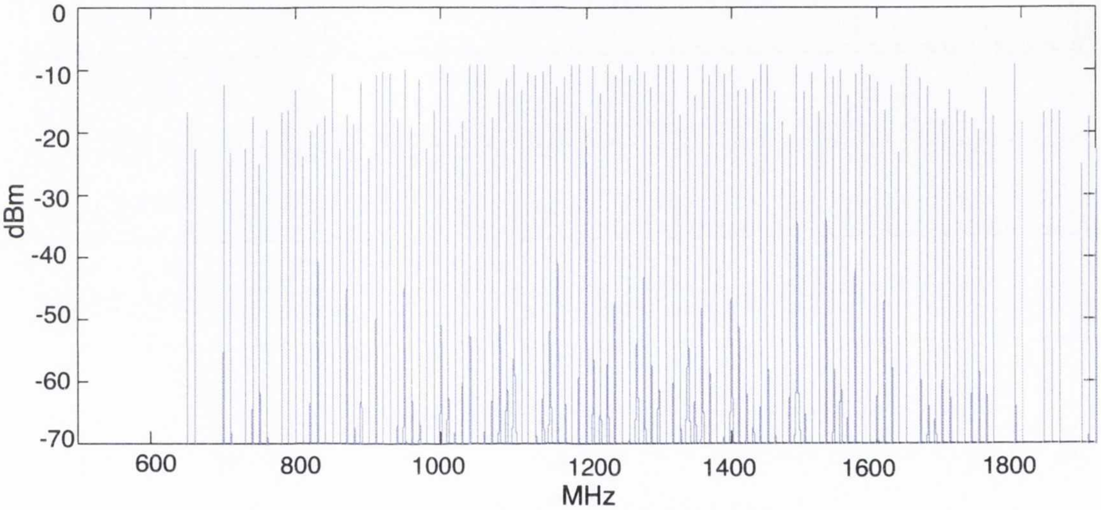

(b) LNA Output Signal Distribution

Figure 3-23 Effect of Low Noise Amplifier Distortion on Spectral Distribution

3.4.2 Spectrum Environments Characterization

The characterization of the spectrum environments that impact front-end linearity is an analog of the spectrum occupancy considerations discussed in the previous section with two important distinctions.

1. The linearity performance is driven by the high-energy region of the spectral energy distribution.
2. The technology used in front-end pre-selector filters typically has pass-band width that is proportional to frequency, rather than fixed or absolute, as in IF or digital filters.

In a demodulator, it is reasonable to treat the response bandwidth as independent of frequency[1]. In a front-end filter, the bandwidth is both significantly larger (generally), and

---

[1] For example, a 10 MHz Intermediate Frequency stage response range is constant whether the receiver is tuned to 300 MHz or 3 GHz.

it is typically proportional to the tuned frequency, assuming a fixed filter Q (or pole count). A filter with a given complexity might have a bandwidth of 100 MHz at 1 GHz, and one of similar complexity would be expected to have a bandwidth of 200 MHz if designed to operate at 2 GHz. For that reason, front-end filter bandwidth is treated as a constant proportion of the center frequency, rather than an absolute bandwidth. Whereas the signaling bandwidth (b_0) discussion was focused on protecting the digital demodulation process from in channel noise, the pre-selector bandwidth ratio (BW) concern is for the operation of the linear analog stages of the receiver. A histogram of spectral power is shown in Figure 3-24, for a 25 kHz bin bandwidth.

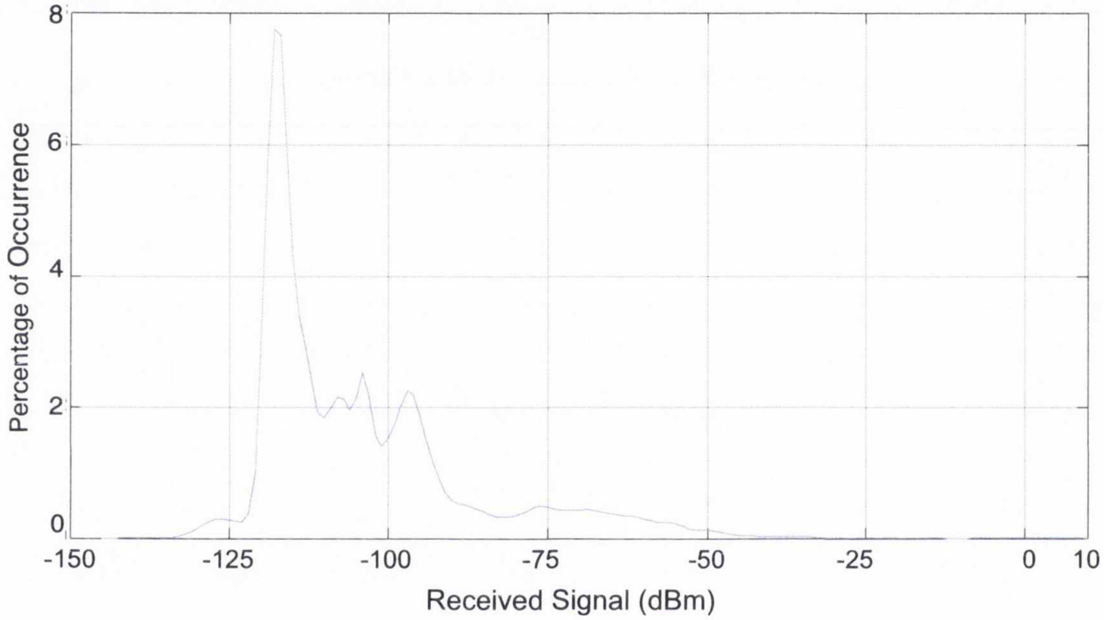

Figure 3-24 Histogram of Spectral Power Levels for Chicago Collection (Linear)

Although this chart would appear to indicate that low energy dominates the spectrum, the same sample in a log expression demonstrates that a large amount of energy is received in a few sparse, but highly energetic frequencies. More than 1 in 10^5 individual frequencies (25 kHz resolution channels) have power in excess of 1 dBm! The same data is depicted in a log format in Figure 3-25. Additionally, a practical filtering system would be very likely to allow multiple signals to enter the front-end, so the energy would be additive in driving the front-end into non-linear behavior.

Equally concerning is the aggregation of energy in closely spaced frequencies. Even with reasonable front-end filter performance levels, energy from multiple channels would be aggregated when presented to the LNA and first mixer stage. The correlation of the energy in individual frequencies is critically important to understanding the LNA's environment.

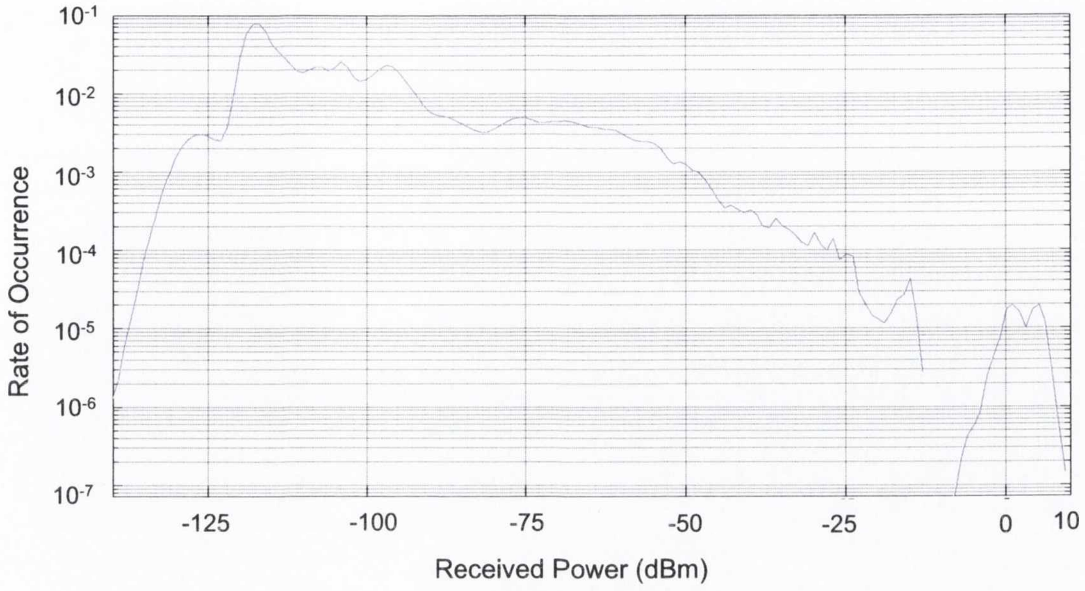

Figure 3-25 Histogram of Spectral Power Levels for Chicago Collection (Log)

There are several inter-modulation products to consider. The second-order inter-modulation creates mixing products that are at the sum and difference frequencies. For a receiver with pre-selector selectivity of less than an octave, these artifacts must fall outside the octave band of interest ($f_1 \pm f_2$, when f_1 and f_2 are within the pre-selector filter pass-band). Signals that would have mixing products within the band of interest are outside of the pass-band of the input filter, and thus, are not relevant to a receiver with less than one octave filtering [130]-[133]. Less amenable to filtering strategies, and of particular concern to a less than octave filtered radio are the third-order intercept products, some of whose mixing products fall into the original tuning range of the front-end filters ($f_1 \pm f_2 \pm f_1$), creating products mixing the input frequencies with the differences between the input frequencies.

The intermodulation response is measured by an intercept point, which is the input energy that causes the linear response to be equal to the intermodulation (of that order) response. IIP2 is the second order, and IIP3 is the third. Typical IIP3 values range from -12 dBm for lower cost, low power consumption devices, to many dBm for high cost, higher performance devices.

An example of total energy to the front-end LNA stage for various absolute bandwidths is shown in Figure 3-26 for the Chicago sample[1].

---

[1] This chart depicts the bandwidth in absolute terms in order that the linearity of the effects is clear. The rest of this section will consider bandwidth as a measure proportional to the center frequency.

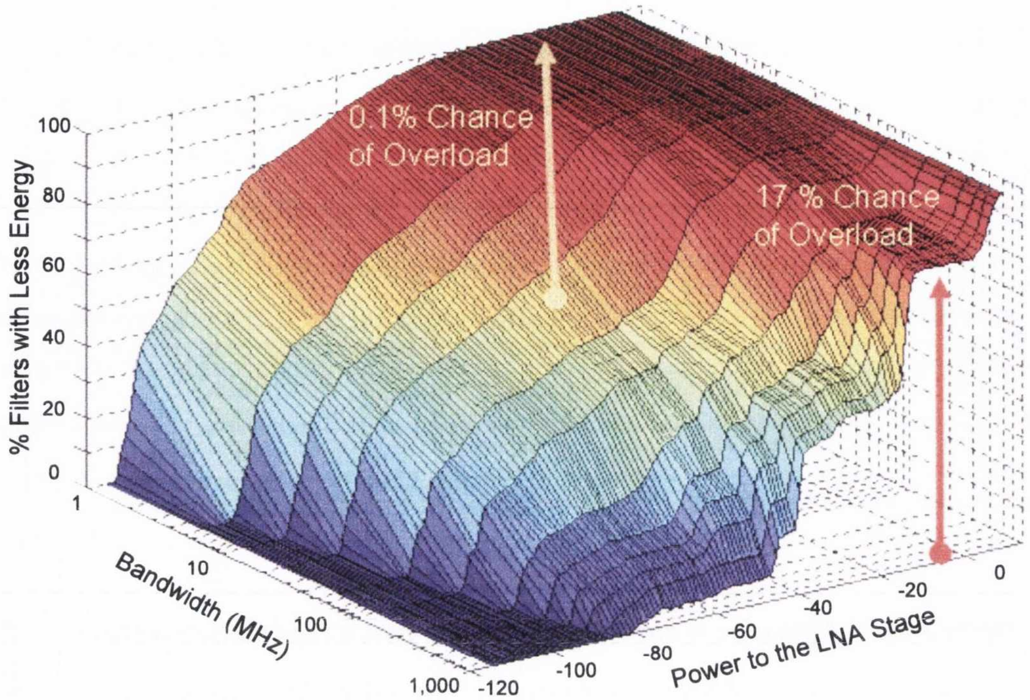

Figure 3-26 Total Spectral Power as Function of Filter Bandwidth for Chicago Collection

In the chart above, it is worth noting the intercept of the power surface with the -5 dBm contour, typically the higher end of the IIP3 for medium-cost consumer devices. In an inflexible frequency assignment process, such as manually assigned spectrum, any situation in which this energy level was exceeded, or even approached would have significant link margin degradation for operation on any frequency within the filter band. This operating point is thus a critical measure of RF component performance. In this case, the -5 dBm point is only reached 0.1% of the time when the filter bandwidth is 1 MHz, but has a 17% chance of being above the -5 dBm level when the bandwidth is 1 GHz. As will be shown later, performance is highly compromised well before this IIP3 level is approached.

An important note is that receiver dynamic range is not just a product of good (or alternatively, bad) design. There is an inherent relationship between the linearity achieved in an amplifier stage, and the energy required to power it for a given gain level. Additionally, the required energy is not limited to that required to power the front-end LNA alone, but also has a proportional effect on energy required in other stages, such as the local oscillator and mixer stages. Higher linearity drives higher energy consumption, which in turn drives energy storage, thermal load and dissipation, substrate area, and ultimately, cost.

A number of strategies have been employed to mitigate the receiver desensitization and non-linear effects of the radio environment. One commonly applied approach is Automatic Gain Control (AGC). This adjusts the receiver gain in one or more stages to remain within the dynamic range of the circuit. AGC to mitigate a stronger than necessary desired signal has negligible effect on link operation, as it is only required in the situation in which excess signal is present. The receiver dynamic range is generally far in excess of the flat BER portion of the Shannon curve (E_b/N_o), so reducing the front-end gain is acceptable, when the energy controlling the AGC level is the desired signal. However, if AGC operation is required due to the energy of an adjacent in-band signal, then the loss of gain raises the effective receiver noise floor, and therefore reduces link performance. AGC is an effective technique for in-channel energy, but compromises link performance if required to reflect energy of adjacent signals[1].

3.4.3 Non-Linearity Effects on Spectrum Sensing and Receiver Operation

The spectrum measurements make it possible to examine and quantify the opportunities for effective dynamic range management by a cognitive radio. The key elements of this technique are selection of filter frequency, and potentially, filter width and its associated insertion loss. Filter width is generally expressed as a filter Q (or Quality Factor), which describes the equivalent Resistor-Capacitor-Inductance (RCL) circuit. One important characteristic of filters is the shape of the filter's skirts, or how much attenuation is provided for signals not within the pass-band of the filter.

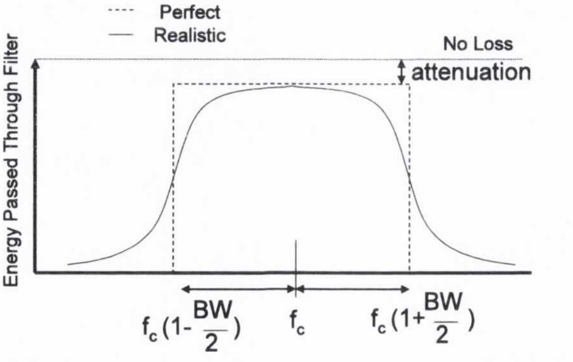

Figure 3-27 ***Perfect Filter Equivalent compared to a Realistic Filters***

To be able to characterize all of these effects from the perspective of a cognitive radio, the filter is modeled as perfect, or ideal, as shown in Figure 3-27. Since the major consideration is energy passing through the filter, the non-ideal performance of the filter is characterized in terms of the equivalent energy passed through an "ideal" filter, given the same central frequency attenuation. This characterization is an appropriate one, since the cognitive radio will make decisions based on environments as they are measured by the

---

[1] When AGC response causes the demodulator to lose performance in demodulating the signal that, in the absence of an in-band interferer it could acceptably demodulate, it is referred to as "desensitization". Desensitization from the intended signal itself is not a concern, since the receiver operates in a high SNR.

device, and the overload behavior is driven primarily by energy, not spectral distribution. This characterization is considered to be the "noise bandwidth" of the filter, and closely approximates the 3 dB bandwidth points of a realistic filter with more than two poles. The filter bandwidth will be characterized as *BW*, which is a fixed proportion of the center frequency (f_c). This representation is relatively proportional to filter complexity, as measured in filter poles[1]. Considerable design evolution has gone into controlling the relative bandwidth of the pass-band, and the shape of the filter skirts, but this dissertation will consider filter response as a percentage of center frequency. This model can be applied to any set of spectral occupancy data.

Filter bandwidth is important for two reasons. The total energy into the front-end is clearly impacted by the width of the filter. If the signals in a segment of spectrum are typically spaced less than the bandwidth of the filter, it is reasonable to assume that the total energy is proportional to the width of the filter. If the energy is dominated by a single or small set of signals, the introduction of narrow filters provides the cognitive radio with more choices (statistically independent trials) that can enable it to select a frequency in a band in which strong signals are not present. The spectrum sample data provides regions where both environments are evident, as annotated in Figure 3-28 [113].

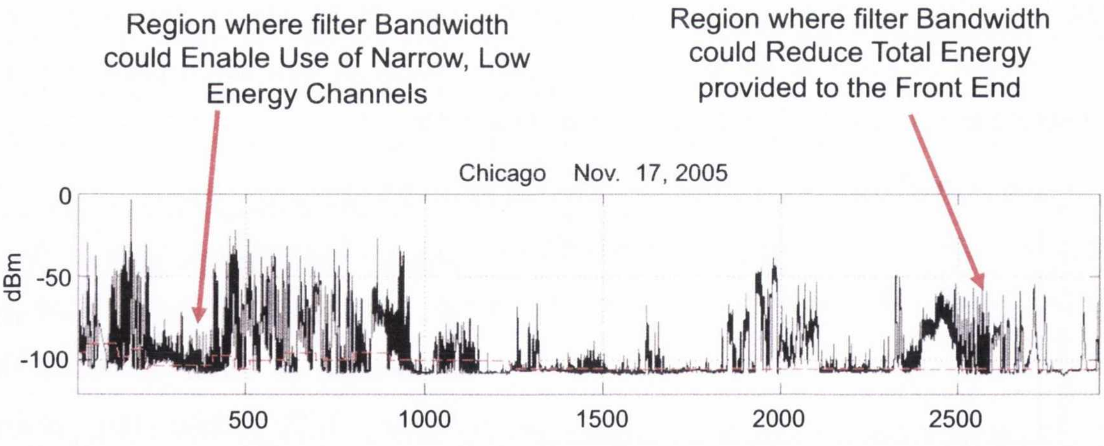

Figure 3-28 Different Environments Lead to Different Impacts of Filter Bandwidth

Filters on the order of an octave have relatively little ability to isolate low energy bands, while filters of bandwidths better than 0.20 show ability to isolate strong signals and provide relatively low energy candidate bands of operation. Figure 3-29 illustrates the

---

[1] By contrast, measuring it in absolute bandwidth would mean that the filter's complexity (as measured by poles) would have to vary with center frequency, as a filter at one octave higher in frequency would have to be twice as selective to maintain a constant absolute bandwidth. The selected measure is appropriate to examine the relative trades between performance and complexity in a frequency independent framework, without constraints of specific filter implementation technique.

energy passed by the filter for a range of bandwidths and center frequencies, using the Chicago spectrum occupancy data.

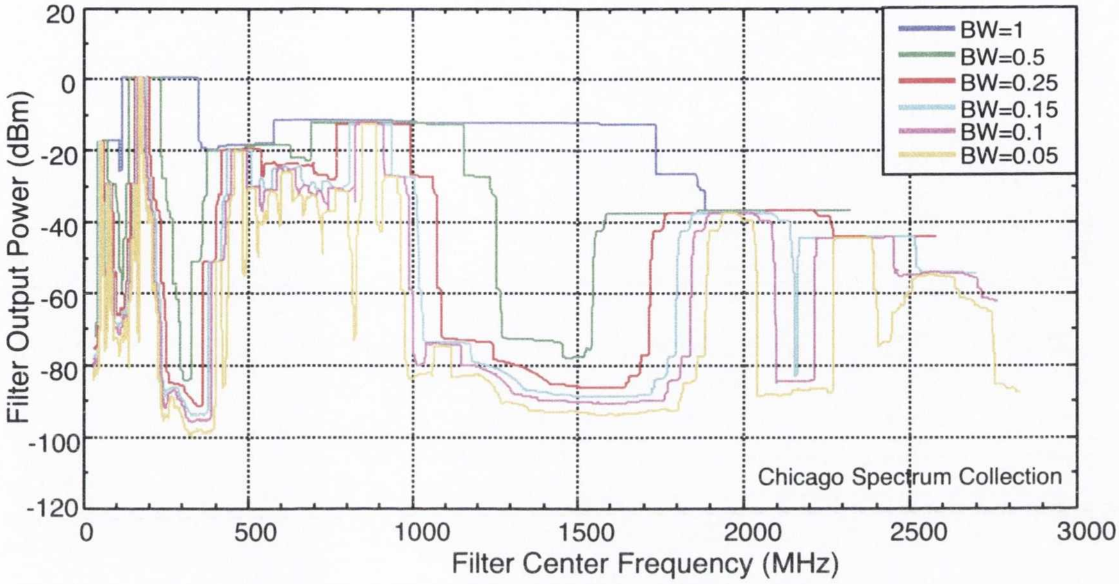

Figure 3-29 Filtered Spectrum Power Data

Another perspective of this distribution is by the distribution of energy levels. This is shown in Figure 3-30.

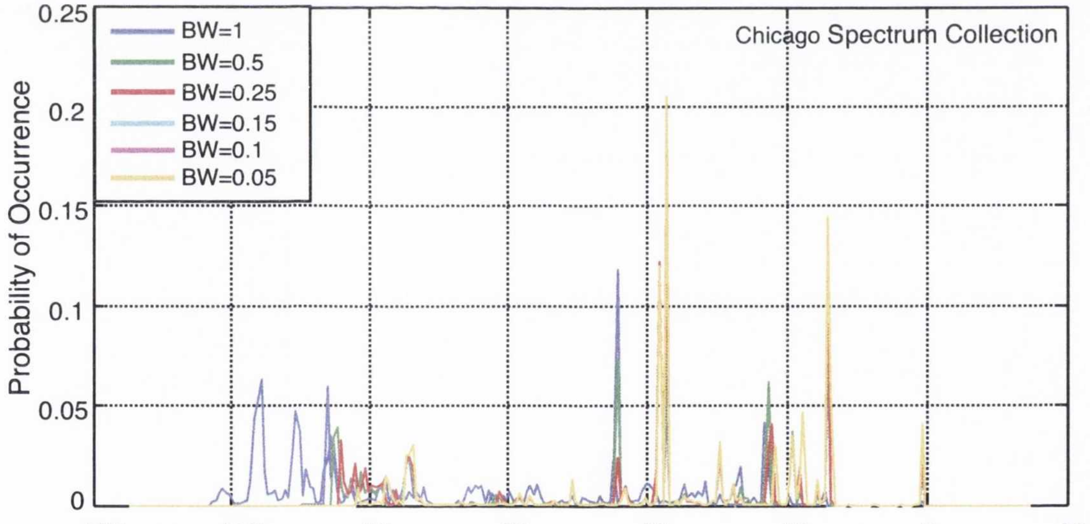

Figure 3-30 Illustrative Distribution for Bandwidth Constrained Spectral Power

Figure 3-31 illustrates the CDF for the Chicago spectrum for a range of filter bandwidth ratios. Similar to the analysis of spectrum occupancy, this distribution is clearly skewed to the low energy end of the distribution at small bandwidths, and to the high-energy end for larger bandwidths. The distribution reflects the correlating effect of different usage patterns caused by both regulation and the inherent technology and physical constraints.

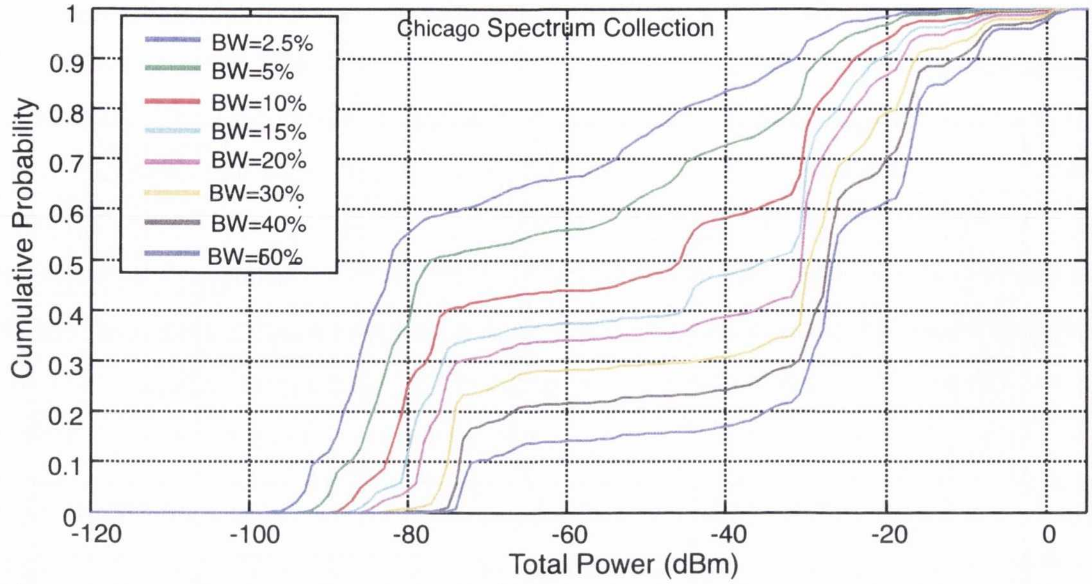

Figure 3-31 ***Illustrative Cumulative Distribution for Bandwidth Constrained Spectral Power***

Figure 3-32 provides the 20% bandwidth *(BW* = 0.2) energy for a number of the spectrum collections.

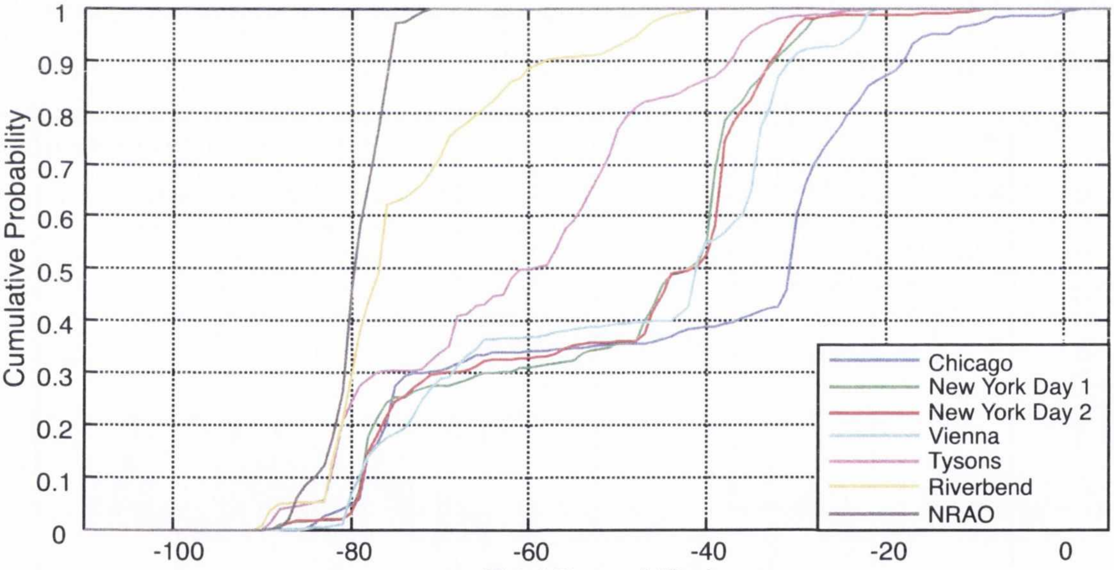

Figure 3-32 ***Power Distributions for a Number of Locations with BW = 0.2 Pre-selector Bandwidth***

The characterization of front-end energy distribution will follow a similar structure to that used to characterize spectrum occupancy of signal bandwidths, except considering both the much wider bandwidth of front-end pre-selectors, and the fact that their bandwidth is generally proportional to the operating frequency.

To determine the CDF, each distribution is characterized as previously shown in Figure 3-6. For the front-end, the energy is within the range of $FE_{\min}$ to $FE_{\max}$. The *FE* Energy

CDF is non-zero above the value of FE_{min}, and is unity at all values above FE_{max}. The FE_{min} threshold levels for each spectrum sample are shown in Figure 3-33.

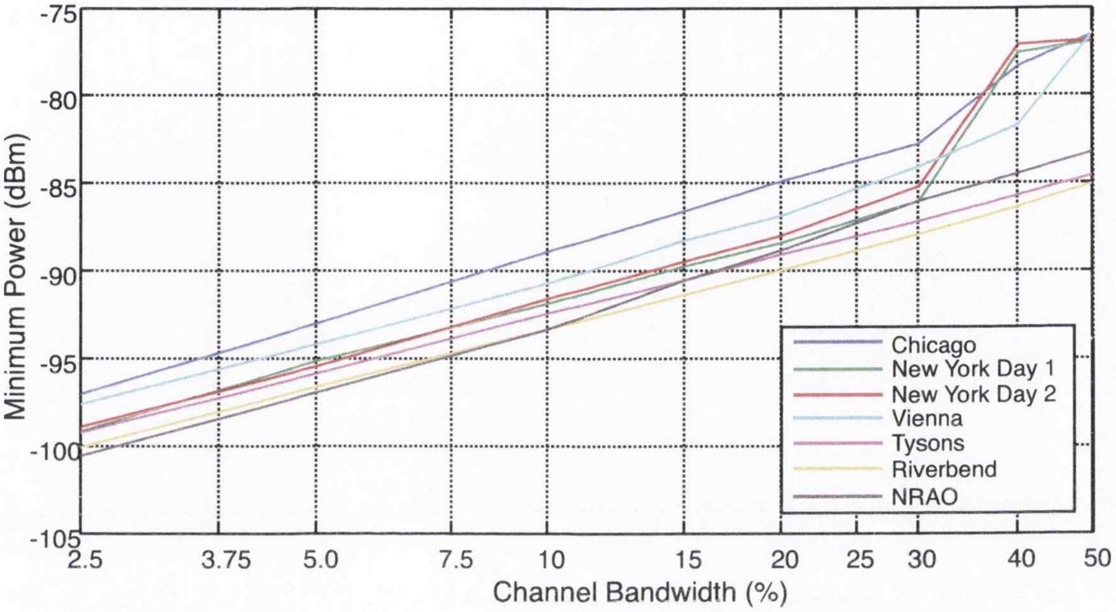

Figure 3-33 FE_{min} Noise Levels for Front-end Bandwidths in Spectrum Samples

As would be expected, the lowest values of channel energy are decorrelated (noise dominated) and are essentially linear with the bandwidth, reflecting that they represent noise bandwidth. At very wide bandwidths, the correlated nature of spectrum usage causes a disproportional increase in total energy in the dense urban environments (Chicago, New York, Vienna), but not for the more rural ones. This behavior implies that at wide bandwidths, all of the urban samples have at least one strong emitter within the filter band-pass.

The corresponding upper energy limits (FE_{max}) are shown in Figure 3-34. Note that several individual emitters dominate these values in each spectral range and sample, and therefore, they are not proportional to bandwidth, but tend to reflect the peak power emitter in each environment. It is clear that the dense urban environments are quite distinct from the rural ones, and again differences between the urban measurements are not fundamental. As in the signaling bandwidth analysis, the maximum energy (FE_{max}) is clearly dominated by a few strong emitters, and so is almost insensitive to the bandwidth. The low energy cases exhibit some sensitivity, but in these environments, front-end energy is typically not a significant concern in any case.

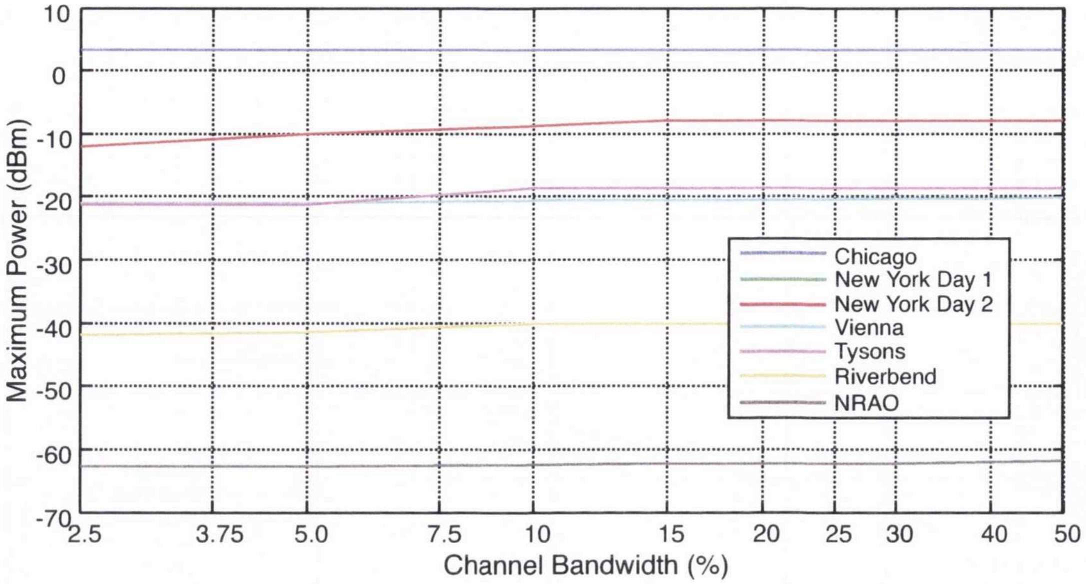

Figure 3-34 ***FE_{max} Maximum Narrowband Energy Levels***

The distribution between the two extreme energy points (the range between the $FE_{\min}$ and $FE_{\max}$ values) can be characterized by the statistical mean and variance of the power. Figure 3-35 illustrates the mean value of this distribution. As in the signaling bandwidth analysis, a small (less than 0.5) value indicates a skew towards the low energy signals, 0.5 is a symmetric distribution, and values above 0.5 indicate that the distribution is skewed towards higher energy levels. When large bandwidths are presented to pre-selector filters, the aggregate energy skews towards the high-energy distribution. The slight non-monotonic behavior at wide bandwidths is due to increases in the FE_{min} discussed previously.

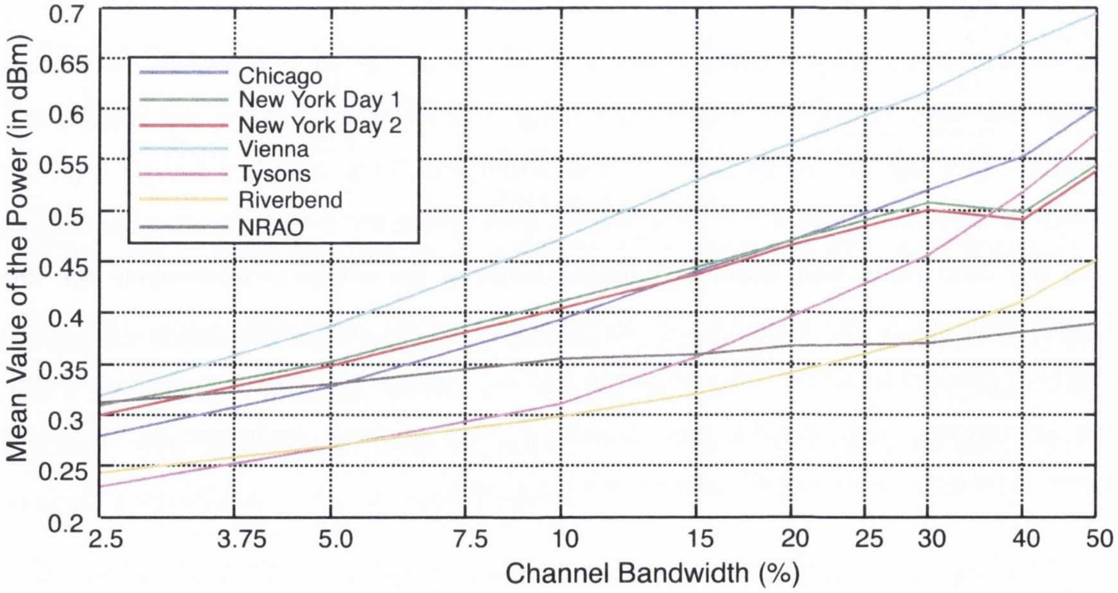

Figure 3-35 ***Mean of Front-end Power Distribution Between FE_{min} and FE_{max}***

The corresponding variance of the sample distributions is shown in Figure 3-36.

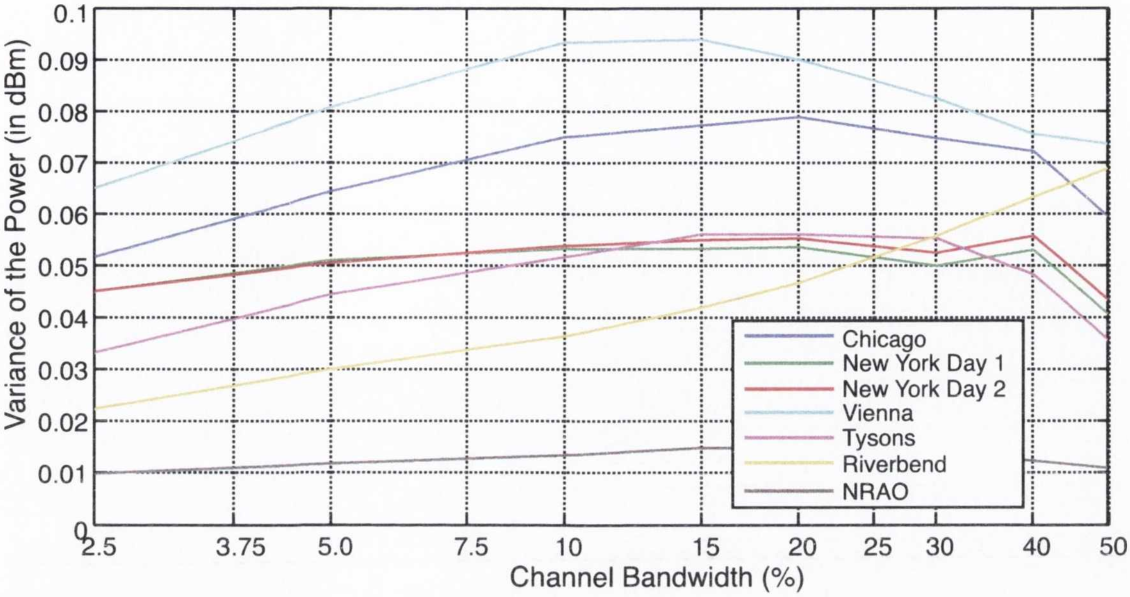

Figure 3-36 Front-end Power Variance of Spectrum Samples

As the mean increases past the mid-point (symmetric point), the variance must decrease accordingly. The beta distribution α and β parameters are derivable directly from the mean and variance, and are shown in Figure 3-37 and Figure 3-38 respectively.

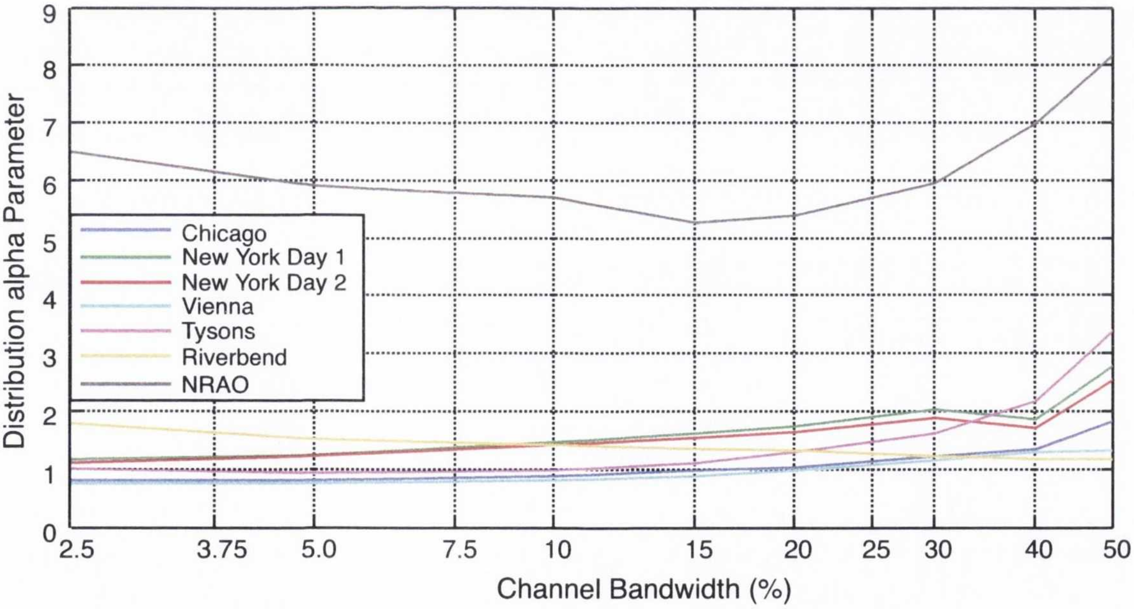

Figure 3-37 Front-end Power Beta Distribution α Value for Spectrum Samples

Again, the NRAO values reflect the correlation and scarcity of the emitters in the environment. However, these measures are quite characteristic and stable for the dense environments that are of the most concern as potential sources of intermodulation overload. The alpha has a very limited range, and the beta converges with high spectral density.

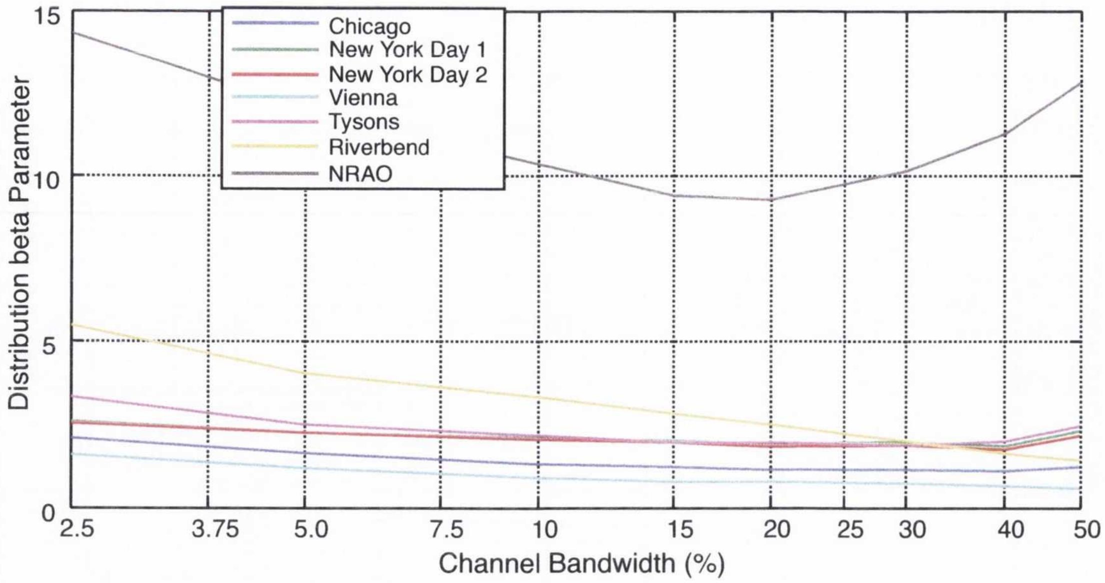

Figure 3-38 Front-end Power Beta Distribution β Value for Spectrum Samples

The values of the minimum and maximum energy, and the distribution α and β parameters are usable directly from these graphs to characterize environments and understand performance sensitivities. It is also useful to develop a closed-form expression of the parameters of the distributions for interpolation. Each of the parameters is modeled as a polynomial equation of the pre-selector bandwidth ratio (*BW*). These equations are shown below.

$$FE_{\min} = k_{\mathrm{FEmin2}} \log_{10}(BW)^2 + k_{\mathrm{FEmin1}} \log_{10}(BW) + k_{\mathrm{FEmin0}} \text{ (in dBm)} \tag{3-8}$$

$$FE_{\max} = k_{\mathrm{FEmax2}} \log_{10}(BW)^2 + k_{\mathrm{FEmax1}} \log_{10}(BW) + k_{\mathrm{FEmax0}} \text{ (in dBm)} \tag{3-9}$$

$$FE_{\mathrm{alpha}} = k_{\mathrm{Fe}\alpha 2} \log_{10}(BW)^2 + k_{\mathrm{FE}\alpha 1} \log_{10}(BW) + k_{\mathrm{FE}\alpha 0} \tag{3-10}$$

$$FE_{\mathrm{beta}} = \mathrm{k}_{\mathrm{So}\beta 2} \log_{10} BW)^2 + k_{\mathrm{FE}\beta 1} \log_{10}(BW) + k_{\mathrm{FeE0}} \tag{3-11}$$

The estimators' polynomial coefficients (mean square fit) for each sample set are shown in Table 3-6 and Table 3-7.

Table 3-6 Empirically Derived Front-end Distribution Constants

	CDF *FE* α			**CDF *FE* β**		
Samples	$k_{\mathrm{FE}\alpha 2}$	$k_{\mathrm{FE}\alpha 1}$	$k_{\mathrm{FE}\alpha 0}$	$k_{\mathrm{FE}\beta 2}$	$k_{\mathrm{FE}\beta 1}$	$k_{\mathrm{FE}\beta 0}$
Chicago	0.798	0.345	3.21	2.02	-6.3	9.61
NY Day 1	1.18	1.94	1.74	2.6	-5.36	9.47
NY Day 2	1.15	1.79	1.28	2.6	-5.59	9.36
Vienna	0.659	1.83	-0.915	1.5	-5.15	6.96
Tysons	1.04	-1.9	12.8	3.22	-10.2	17.7
Riverbend	1.73	-3	3.86	5.31	-18.6	22.4
NRAO	6.52	-11.3	29.9	14.1	-37.1	71.6

Table 3-7 Empirically Derived Front-end Energy Constants

	Signal Floor (FE_{min})			Signal Ceiling (FE_{max})		
Samples	k_{FEmin2}	k_{FEmin2}	k_{FEmin2}	k_{FEmax2}	k_{FEmax1}	k_{FEmax0}
Chicago	125	-343	368	0.118	-0.343	0.319
NY Day 1	78.5	-116	106	56.9	-190	194
NY Day 2	77.5	-91.4	61.5	57	-190	194
Vienna	124	-363	421	8.72	-28.9	32.5
Tysons	109	-288	274	47.4	-161	165
Riverbend	108	-282	270	27.5	-94.4	98.2
NRAO	119	-285	251	4.95	-14.7	16.7

To obtain an estimate of how effectively these polynomial expressions reflect the parameters of the distributions, the measured and closed-form estimates of minimum and maximum energy are illustrated in Figure 3-39. This chart illustrates the fit of the closed-form estimates of the minimum and maximum energy polynomial coefficients, with the values derived from the measurements for the Chicago sample. Similarly, the fit for the alpha and beta parameters are provided in Figure 3-40.

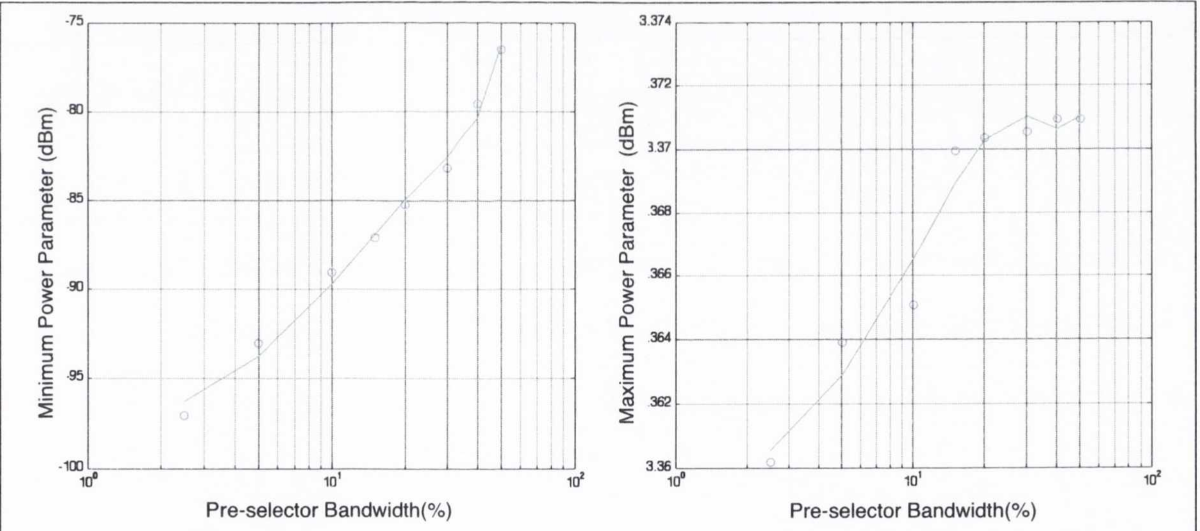

FE_{min} Parameter	**FE_{max} Parameter**

Figure 3-39 Empirical and Closed-form Distribution for Chicago Minimum and Maximum Front-end Power

Note that while the polynomial estimates are not "perfect", the variations are small compared to the differences between various spectrum collections, and thus are within the anticipated environmental variation.

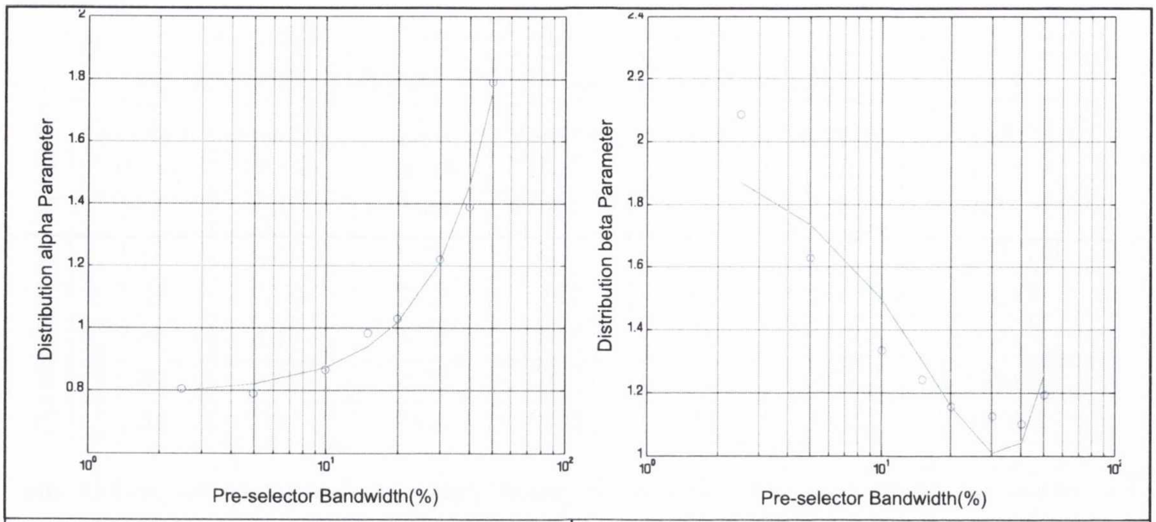

***FE* Beta Distribution α Parameter**	***FE* Beta Distribution β Parameter**

Figure 3-40 Empirical and Closed-form Distribution Parameters for Chicago Distribution Parameters for Front-end Power

The CDF of the spectral energy (probability that energy is between FE_{min} to FE_{actual}) is:

$$\mathbb{P}(FE_{actual}, BW) = \frac{\beta_x(FE_\alpha, FE_\beta)}{\beta(FE_\alpha, FE_\beta)}$$

where:

$$x = \frac{FE_{actual} - FE_{\min}(BW)}{FE_{\max}(BW) - FE_{\min}(BW)} \tag{3-12}$$

FE_{actual} is the Front-end energy whose probability is being determined

To obtain a sense of the accuracy of the resulting CDF, Figure 3-41 provides the measured and estimated occupancy estimates for the Chicago front-end energy distributions.

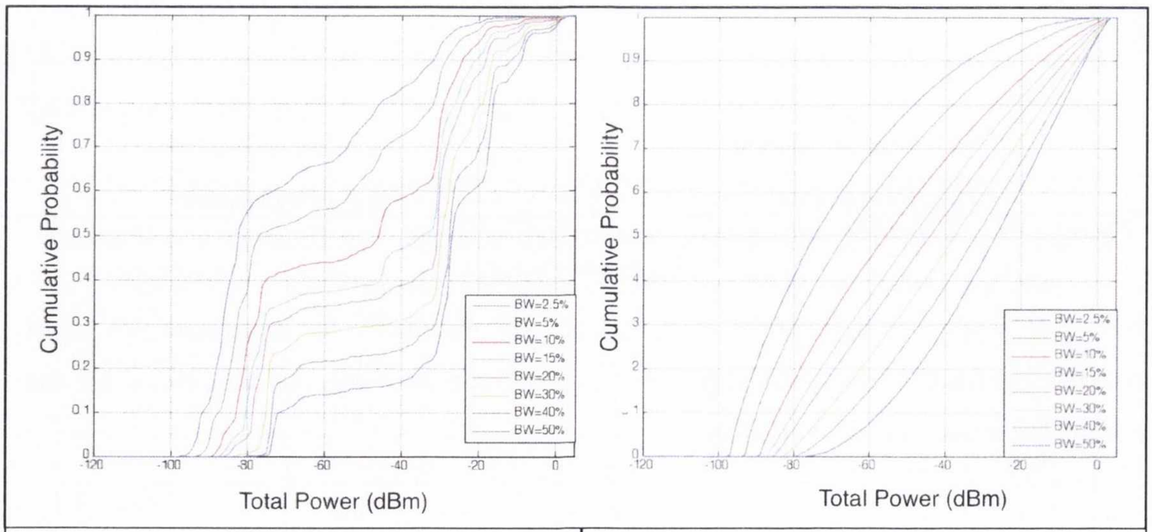

Measured	**Estimated**

Figure 3-41 Comparison of Empirical and Estimated Front-end Power Distributions for the Chicago Spectrum Sample

A summary of the front-end energy distribution CDF estimates for a representative fixed pre-selector bandwidth (20%) for all spectrum samples is shown in Figure 3-42.

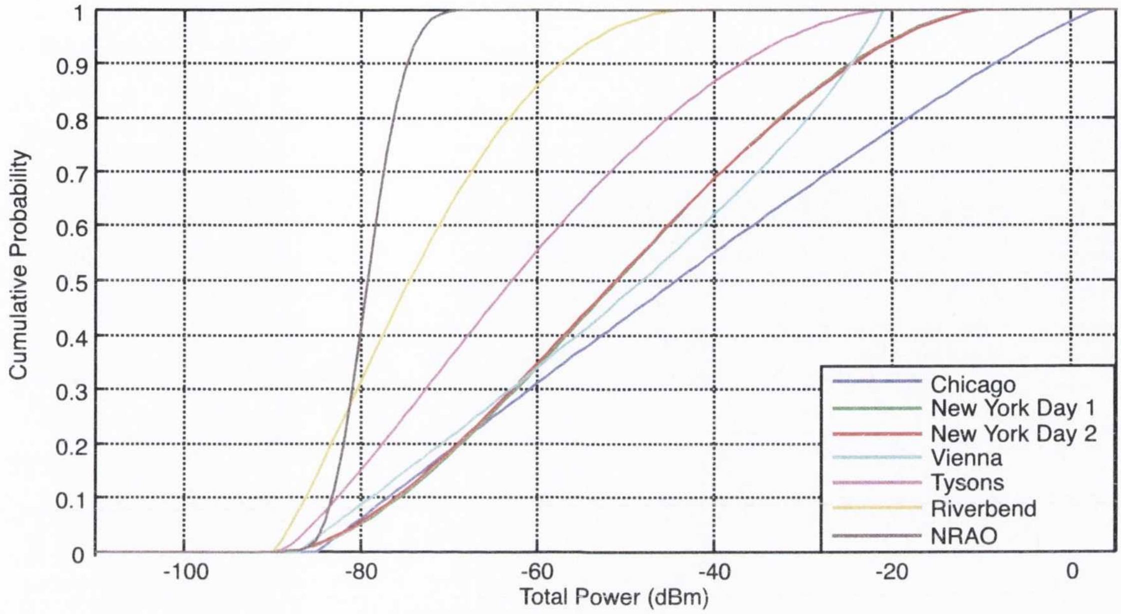

Figure 3-42 Parametric Distributions of Front-end Power for all Spectrum Samples

The benefits of this expression are twofold. First, it provides a convenient and tractable composite of over 7×10^8 individual data points. Second, and most important, it provides a closed-form framework in which assertions about cognitive radio performance can be more globally proven over a class of environments, rather than demonstrated in a single environment.

3.5 Monotonic Classifier of Spectrum Characteristics

The spectrum distribution enables closed-form characterization of a specific spectrum environment. This section extends this structure to enable synthesis and characterization of generalized environments. These environments can be broadly classified in terms of their signal density and the dynamic range of signal energy. While these two measures are related, they warrant independent expression. For example, an extremely high antenna might be able to detect a large number of weak signals, and therefore report dense spectrum usage, although none of the signals might be very strong. Conversely, a sensor in the middle of an urban environment may be shielded from most of the signal density, but the signals it did perceive would be close range with direct paths, and thus might be high energy.

These two measures will be referred to as Density Index ($I_{Density}$) and Intensity Index (I_{Energy}). Although these indexes could be unit-less, a starting point is to associate them with measurable phenomena. The Density Index will be the difference (in dB) in the minimum energy level and the median energy level in one of the smaller channel bandwidth investigated, which is arbitrarily selected to be the 1 MHz bandwidth measurement set. This index reflects the relative density of usage; i.e., how many signals are present to account for 50% of the energy distribution. In a very low density environment, few signals are present, and only a few dB separate the noise from the median; for a dense environment, a large signal spread is present. The Intensity Index ($I_{Intensity}$) reflects the dynamic range (in dB) between the strongest and weakest signal in a small bandwidth (25 kHz).

The Indexes are defined as follows:

I_{Density} = Mean Signal level of the median energy per channel above the noise floor (for a 1 MHz channel)

$I_{\text{Intensity}}$ = Range from the strongest energy to the noise floor measured at the smallest bandwidth (25 kHz)

The two days of the New York City collection are tightly clustered, as would be expected. Figure 3-43 illustrates the clustering of the spectrum indices for the spectrum collections.

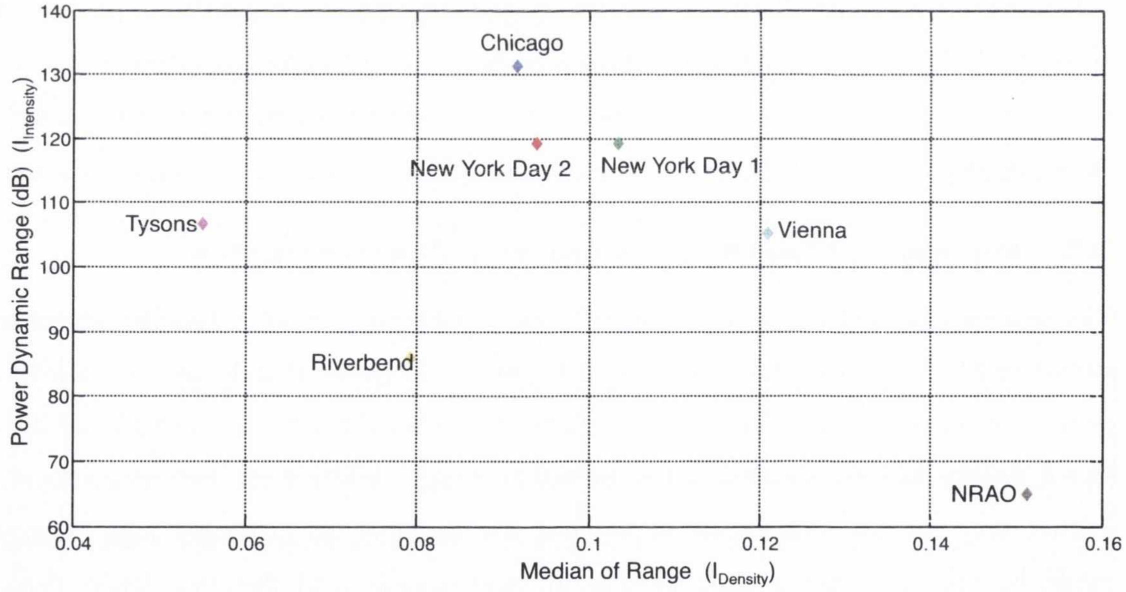

Figure 3-43 Spectrum Characterization Indices

3.5.1 Spectrum Occupancy Classification

The use of a spectrum index requires that the estimating functions of Equations (3-1) thru (3-4) be extended to reflect these indices. The determination of the high and low energy limits, and the alpha and beta parameters of the distribution are estimated using a polynomial with the spectrum environment characteristics and the bandwidth to be considered. The coefficients of the indices and the signal bandwidth are shown in Table 3-8.

Table 3-8 ***Coefficients of Spectrum Signaling Bandwidth Energy Estimator Polynomials***

	Spectral Characteristics				**Signal Bandwidth**		**Constant**
	I_{Density}	I_{Density}^2	$I_{\text{Intensity}}$ *10	$I_{\text{Intensity}}^2$ *100	$\text{Log}_{10}(b_0)$	$\text{Log}_{10}(b_0)^2$	
SO_α	-2.77	20.2	-0.487	0.0213	0.0692	-0.0399	2.8
SO_β	-37.5	134	-1.8	0.0788	0.488	-0.184	13.9
$SO_{\min}$	-60.4	137	-2.12	0.0977	8.59	0.503	-125
$SO_{\max}$	-69.7	189	7.67	0.0997	11.6	-1.5	-131

3.5.2 Spectrum Energy Classification

A similar table of polynomial coefficients can be constructed for the front-end overload energy as a function of the spectrum indices and the pre-selector bandwidth ratio. These coefficients are shown in Table 3-9.

Table 3-9 ***Coefficients of Spectrum Pre-selector Bandwidth Energy Estimator Polynomials***

	Spectral Characteristics				**Filter Bandwidth**		**Constant**
	I_{Density}	I_{Density}^2	$I_{\text{Intensity}}$ *10	$I_{\text{Intensity}}^2$ *100	$\text{Log}_{10}(BW)$	$\text{Log}_{10}(BW)^2$	
FE_α	-26.3	153	0.0604	-0.00305	0.638	0.247	1.26
FE_β	-84.2	396	-1.19	0.0521	0.272	0.372	11.4
$FE_{\min}$	185	-555	3.37	-0.115	23.2	4.79	-108
$FE_{\max}$	1.49	313	22.2	-0.586	-1.31	-1.55	-189

Using these coefficients, it is possible to develop spectral distributions for any plausible spectrum environment, even in the absence of measured signals. For example, one could hypothesize an urban environment in which mobile nodes where immersed in close proximity to other nodes, but would also be blocked from many other sources by buildings and terrain. In that case, it is reasonable to assume that the spectrum would be more energetic than the rooftop samples provided in this analysis, but also, less dense due to the limited visibility to the mass of devices inhabiting the city. Considering the Cartesian classification of Figure 3-43, the $I_{Intensity}$ measure could be adjusted to reflect the distance and power of the devices that would be in close proximity to the cognitive radio, the $I_{Density}$ measure decremented to reflect the lowered visibility of other devices, and a set of distributions generated. More generally, performance can be proved over a region of distribution characteristics, and then mapped to the spectral environments for which this proof was applicable.

3.6 Summary

Although individual spectrum distributions have a large number of variations, the basic distributions are consistent, and can be parametrically described using continuous closed-form distributions, whose parameters are derivable from spectral measurements. This analysis provides a mechanism to:

1. Simulate a wider range of spectrum environments than can be sampled and analyzed.
2. Perform analysis of radio performance without researchers having large databases of environments.
3. Make provable assertions about cognitive radio performance in large regions of any potential environment.

The number of spectrum sample sets currently available is quite limited, so the approach provided here will significantly benefit from additional spectrum samples, primarily to remove the inherent correlation present in the small sample base. Additionally, the inclusion of additional environments will enable extension of this approach to deal with the non-independent, correlated nature of the spectrum environment, which could not be performed with the limited experimental base currently available. Of particular importance is additional characterization of "outlier" environments. It is likely that the NRAO sample is as sparse an environment as is present in normally populated regions, but at the other end of the scale, more measurement sets are needed to better understand the effects of emitter density in urban and highly co-located environments.

An additional thrust of future spectrum collections should be to address the dimension of spatial and temporal correlation. The degree of geographic and temporal correlation that will be encountered is important if the performance of environmentally adaptive wireless devices are to be fully characterized throughout the range of environments in which they will operate.

Chapter 4 Front-end Linearity Management

4.1 Front-end Linearity Management Objectives

The effects of receiver front-end energy loading are currently a significant factor in the performance of conventional radio devices that inhabit densely populated (RF devices) environments with either discrete high power emitters, or aggregations of emitters collectively resulting in high energy levels. Current manual planning processes have generally introduced an implicit control over the environment to which wireless devices are subjected. Even with effective manual planning, and high performance RF equipment, there are numerous instances of interference that have been caused by non-linearity within RF receivers. This situation was briefly introduced in Section 3.4 for purposes of discussion of spectrum analysis. This chapter will apply results of the Chicago spectrum samples to typify the analysis process, and conclude with analysis of performance across the full range of spectrum environments described in Section 3.5. The methodology can be applied to any spectrum distribution of the form developed in Chapter 3.

This problem is important to the DSA and cognitive radio community for three reasons:

1. One of the primary benefits of DSA is operation in bands that currently are dedicated to other uses. DSA radios therefore will be subject to a much greater range of environments than band-specific and frequency managed products are in conventional spectrum practice.
2. The future deployment of dynamic spectrum access and cognitive radios will make this situation become more stressing, as spectrum density is increased by technologies such as DSA.
3. Provision of high front-end performance is a fundamental cost and energy driver for wireless devices, and is a significant operational impediment. Adaptations that minimize the effects of intermodulation enable significant reductions in required IIP3 performance, and increase reliability beyond what is possible through high power and high cost circuit approaches.

As a minimum, overload will reduce performance by raising the effective noise floor, and often may preclude operation of the device at any capacity level. It has been suggested by Marshall [101] that a cognitive radio can adapt the use of spectrum to avoid situations in which the front-end will be overloaded, or desensitized (such as by AGC behavior)[1].

Note that a radio may have to address more than the environmental effects; duplex operation may generate intermodulation effects that fall into the receive bandwidth, particularly if multiple transceivers are in use [134]. In this treatment, it is assumed that

---

[1] This effect is commonly referred to as adjacent channel, co-site interference, or receiver desensitization.

the device is aware of its channel usage, and therefore, can reflect any additional constraints into the decision process.

If the device is operating with static frequency assignments, there is little effective mitigation of this front-end overload condition, other than high performance receiver front-end performance. On the other hand, if the device were permitted to select its own frequencies, then implementing the following basic principle would appear to offer mitigation:

Frequencies should be selected such that the pre-selector tuning can ensure that the total energy passing through the pre-selector is constrained to no more than a certain ratio of the overload input energy.

The actual spectrum environment is outside the control of the receiver. The primary control the receiver has to mitigate or reduce intermodulation is establishing the pre-selector parameters. The implementation of cognitive radio techniques in this area is specific to the hardware capability and organization, but the following analysis is generally applicable to typical configurations and filter technology. It is not the intention of this chapter to provide an in-depth analysis of filter technology: instead, it is intended to develop the fundamental relationships between the capabilities of the filters; the energy behavior of the LNA; the cognitive radio's algorithms; and the link performance impact in a range of environments.

In analyzing receiver operation, the algorithms will consider the two different receiver bandwidth considerations developed in Chapter 3. The input pre-selector bandwidth (referred to as BW, typically a significant fraction of the carrier frequency) is the constraint on energy entering the front-end, and drives the degree to which the front-end LNA and mixing stages are overloaded by total energy. The signaling bandwidth (b_0, typically a fixed bandwidth that is a small fraction of the pre-selector bandwidth) is the range over which additional noise energy will enter the demodulation process.

The cognitive radio must first be able to determine the relationship between the energy in the signal environment, and the amount of noise that would be generated through the intermodulation of signal products. The standard engineering measurement of intermodulation typically inserts two pure tones, and measures the energy in the intermodulation products, which are therefore also tonal. This situation is of reduced concern to a cognitive radio, since the tonal intermodulation products can be avoided through the DSA algorithms, by the same mechanisms that avoid any other occupied frequency. However, in complex environments, due to mixing of the underlying modulation products, it is likely that the intermodulation products will be present in large numbers, and also be of significant bandwidth. When these factors are present, the products are much less tonal, and approach AWGN noise, composed of many individual intermodulation distortion (IMD) products, with energy falling throughout the band-pass range.

Since the n^{th} intermodulation product has a bandwidth n times the original (for equal bandwidths), the effect of the intermodulation is to broaden the intermodulation product significantly, as well as to create multiple occurrences of intermodulation products. For example, if a given piece of spectrum were occupied by 10 signals that each occupied non-overlapping segments that were 0.6% of the bandwidth, the intermodulation would create 100 signals, each of which occupied 1.8%! Collectively, they would appear as noise to any wideband receiver. The broadening effect can be seen in considering two non-overlapping signal inputs, f_1 and f_2; f_1 from f_{1l} to f_{1h} and f_2 from f_{2l} to f_{2h}. One of the second order intermodulation products would be f_1+f_2, which has the lower frequency of $f_{1l}+f_{2l}$ and an upper one of $f_{1h}+f_{2h}$. The frequency extent of this product is therefore $(f_{1h}+f_{2h}) - (f_{1l}+f_{2l})$, or $(f_{1h}-f_{1l}) + (f_{2h}-f_{2l})$. The second order product bandwidth is the sum of the intermodulating component bandwidths. A similar argument exists for the higher order intermodulation products.

Intermodulation products can fall in a wide range of frequencies, but the ones of most concern to a cognitive radio are those generated by LNA intermodulation that fall within an octave range of the inputs; which is the effect of the third order intermodulation[1]. From Rhode [135], the time domain output of an LNA is:

$$V_{out} = G_{LNA} \cdot V_{in} + \frac{-2G_{LNA}^{\;2}}{IIP2_{Volts}} V_{in}^{2} + \frac{-4G_{LNA}^{\;3}}{IIP3_{Volts}} V_{in}^{3} \qquad (4\text{-}1)$$

Where:

V_{in} is the input voltage
V_{out} is the output voltage
G_{LNA} is the LNA Absolute Gain
$IIP3_{\text{volts}}$ and $\text{IIP2}_{\text{volts}}$ are in volts

The definition of the IIP3 point allows the power of the third order intermodulation products of two equal strength signals (P_{Signal}) to be written as:

---

[1] Any odd order product will fall in this range (5th, 7th…), however the third order dominates the energy and will be the basis for this discussion.

$$P_{IMD3} \propto G_{LNA} \frac{P_{Signal}^3}{IIP3^2} \tag{4-2}$$

Where:

P_{IMD3} is the power of the third order product
P_{signal} is the signal power
IIP3 is the IIP3 value in power (dBm)

This relationship would provide a simple and straightforward method to compute intermodulation effects, except that radios are generally subject to more complex environments, with a large number of unequal amplitude signals present in the spectrum provided to the LNA and subsequent stages. However, it is still instructive to understand the general effect of the spectral distribution of front-end energy.

A fixed amount of front-end energy divided evenly among *n* equal amplitude signals, results in total IMD_3 induced noise that is the individual distortion products of the signals times the number of permutations of the two unique input signals:

$$P_{IMD3} = n(n-1)G_{LNA} \frac{\left(\frac{P_{Signal}}{n}\right)^3}{IIP3^2} \tag{4-3}$$

Where:

n is the number of intermodulating signals

It is clear that for a given amount of input energy, the IMD_3 produced is not constant, but is a factor of the distribution of the energy. The IMD_3 is different when all the energy is concentrated into two signals, compared to when the energy is distributed amongst many signals. This is readily apparent by examining the derivative of total IMD_3 (4-3) against the number of signals, for a constant total input energy.

$$\frac{\partial P_{IMD3}}{\partial n} = \frac{(-n^{-2} + 2n^{-3})G_{LNA}P_{Signal}^3}{IIP3^3} \text{ for } n \geq 3 \tag{4-4}$$

Clearly, total IMD_3 is maximized when only two signals are present, and minimized when the energy is most evenly distributed as *n* increases. The real world operating points between these extremes are a function of the characteristics of the individual environments. For a cognitive radio to assess the operational impact of IMD_3, it is necessary to assess the noise impact of an environment in a computationally straightforward manner. Adaptation mechanisms are dependent on the radio's ability to predict the effect of energy on front-end performance to determine the necessity of adaptations to avoid or mitigate these effects.

Assuming a constellation of n signals of amplitudes $\{a_1, a_2, a_3, \dots a_n\}$ and corresponding frequencies $\{f_1, f_2, f_3, \dots, f_n\}$, the third order IMD product of the mixing of any two of them (a_i, a_j) has amplitude $a_{i,j}$ and frequency $f_{i,j}$:

$$a_{i,j} = G_{LNA} \frac{a_i^2 a_j}{IIP3^2}, \text{ and } a_{i,j} = G_{LNA} \frac{a_j^2 a_i}{IIP3^2} \tag{4-5}$$

$$f_{i,j} = 2 f_i f_j \text{ and } f_{j,i} = 2 f_j - f_i \tag{4-6}$$

The total intermodulation power is given by:

$$P_{IMD3} = \frac{G_{LNA}}{IIP3^2} \sum_{i=1}^{n} \sum_{j=1}^{n} a_i^2 a_j \tag{4-7}$$

An example of the energy output of the pre-selector filter in the Chicago sample, for a range of pre-selector bandwidth ratios (*BW*) is shown in Figure 4-1.

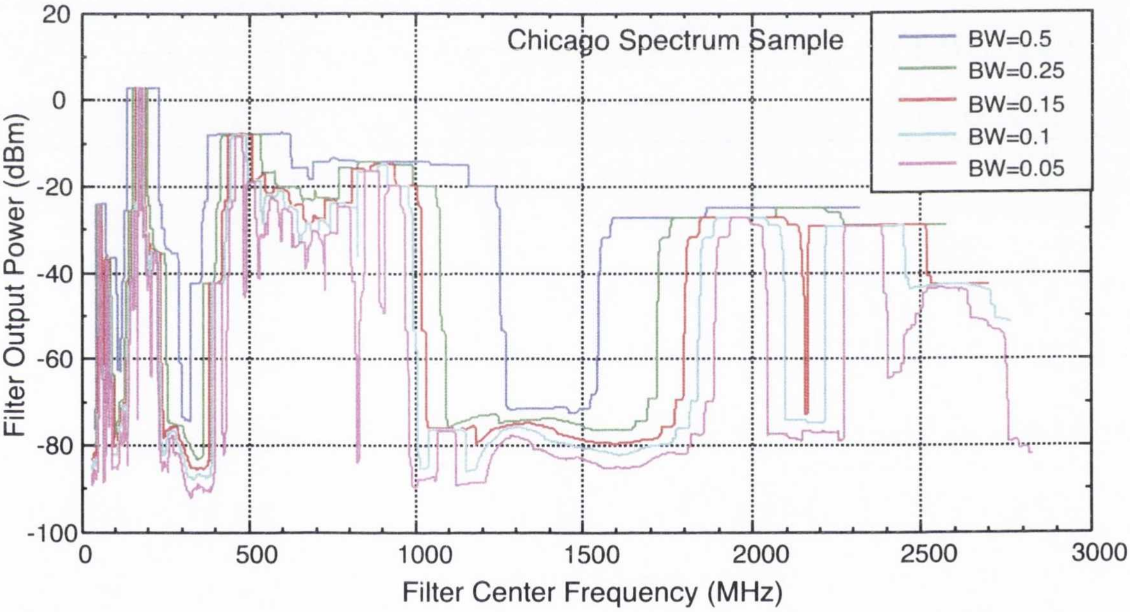

Figure 4-1 ***Typical Filter Output of Pre-selector Filters for a Range of Filter Bandwidth***

Applying (4-7) to the spectrum environments described in Section 3.4 yields the IMD_3 energy distribution shown in Figure 4-2 for a range of filters[1] and frequencies.

For convenience, the energy distributions in this chapter are shown for the equivalent of unity gain, in order to relate the IMD_3 energy products to the equivalent input energy from the antenna ($OIP3 = IIP3$). The bulk of the third order intermodulation energy is distributed over the frequencies between and around the inter-modulating frequencies,

---

[1] The BW parameter interpretation is defined in Chapter 3.

which are the ones passed by the filters. Evenly distributed, this energy appears as additional noise in the demodulator, as if it had arrived through the antenna.

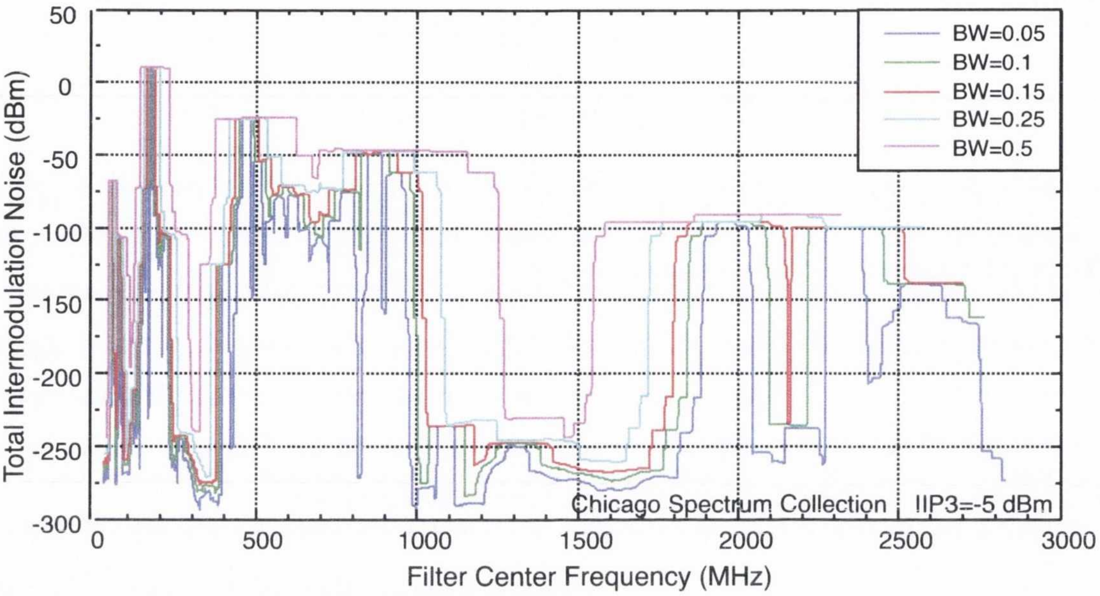

Figure 4-2 IMD3 Power for a Range of Spectrum and Filter Settings

Equation (4-7) can be extended to estimate the mean energy for any given signaling bandwidth b_0 in the proximity of f_c.

$$P_{b0} = \frac{b_0}{f_c BW} \frac{G_{LNA}}{IIP3^2} \left(\sum_{i=1}^{n} \sum_{j=1}^{n} \left(a_i^2 a_j \right) \right) \tag{4-8}$$

Where:

P_{b0} is the power within a signaling bandwidth
f_c is the center frequency of the pre-selector
BW is the bandwidth ratio
b_0 is the signaling bandwidth, and $b_0 < f_c BW$

The distribution of this energy is shown in Figure 4-3 for a 25 kHz wide channel. These values can be scaled to any signaling bandwidth up through the intrinsic bandwidth of the filter. The dotted line reflects a typical receiver noise floor for comparison, and shows that a number of filter bands would have the link noise floor impacted by IMD_3 even for reasonably high values of IIP3.

The energy shown above the -170 dBm/Hz line represents likely significant noise floor elevation in the b_0 pass-band of the demodulator, and would directly reduce link margin. The degradation would likely be so severe that the link would fail completely, since the noise elevation is over 60 dB in some cases. Also, note that narrow bandwidth is not enough to reduce the total intermodulation significantly, since the intermodulation products are due to a small number of excessively strong signals. The reduced filter bandwidth does reduce the extent of frequencies over which receiver operation is so disrupted.

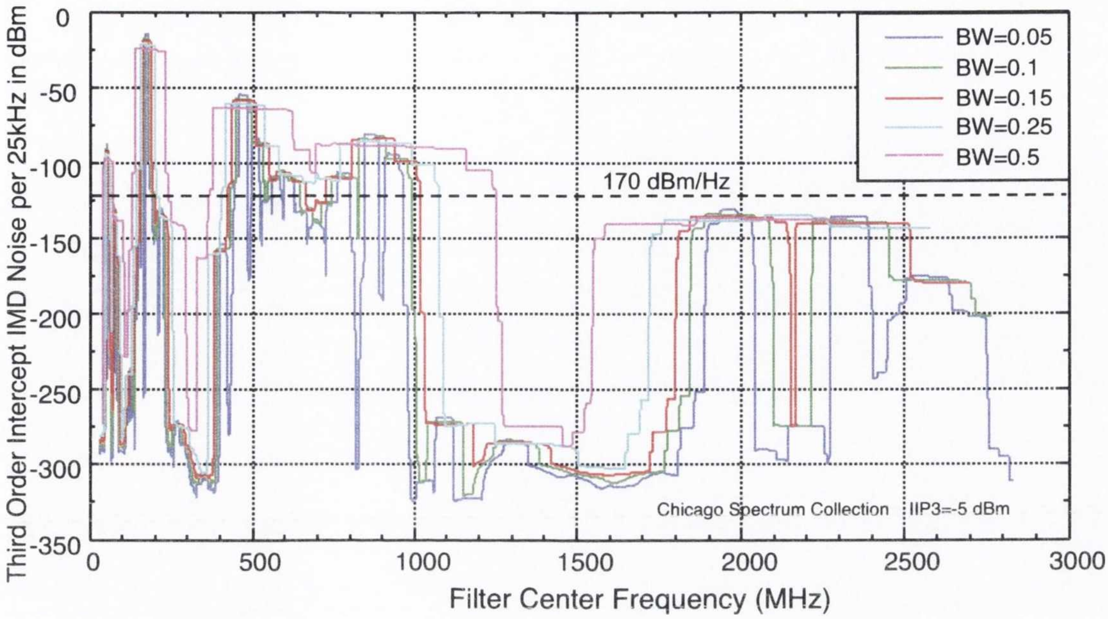

Figure 4-3 IMD3 Noise Power Induced on a 25 kHz Channel

The predictability of the mapping of input energy to the mean of the distribution of IMD3 product is important, as the accuracy of this mapping is a constraint on the ability to utilize total input energy as a criterion for pre-selector selection. Clearly, an estimate based on a directly and easily measured quantity is superior to performing extensive computation and high resolution sensing. Figure 4-4 illustrates the relationship between the input energy and IMD_3 energy for the same set of spectrum measurements and a specific value of IIP3. The dotted red line reflects the pure two-tone response to the same energy.

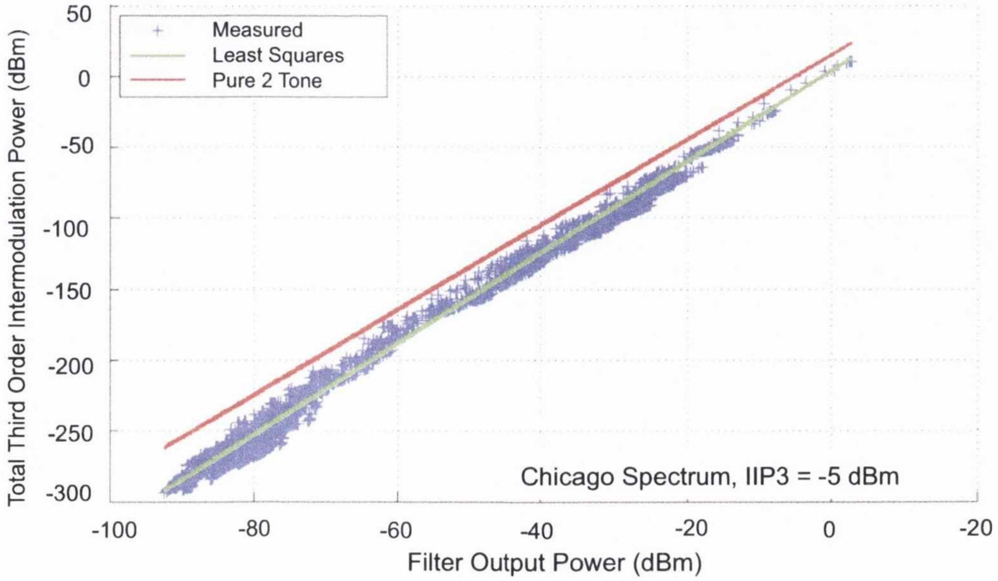

Figure 4-4 IMD3 Power as a Function of LNA Input Energy

The actual IMD3 noise is less than the two tone input response, since the input energy is distributed across multiple signals. At the high-energy end of this curve, the input energy is increasingly dominated by a smaller number of signals, so the response approaches the two tone response. The clusters associated with different IIP3 values are separated by the square of the *IIP3* value difference (2 times the difference in dB).

Applying a least-squares fit to these collections yields a first order polynomial fit to estimate the IMD_3 noise from the input energy:

$$IMD_3 = k_1 P_{\text{in}} - 2\, IIP3 - k_2 \qquad (4\text{-}9)$$

$$k_1 = 3.25,\ k_2 = 11.8$$

where IMD_3, $IIP3$ and P_{in} are in dBm by convention

These coefficients are consistent with the closed-form (which has a first order coefficient (k_1) of 3.0) and indicates an average of approximately 12 dB less intermodulation energy (k_2) than would be expected if all of the input energy was concentrated into the minimum of just two tones, at the point where the input is equal to the IIP3 value. The increased slope (3.25 vs. 3.0) reflects a slight shift in energy distribution of higher energy environments towards more energy concentrated (or correlated) in stronger signals. As energy increases, and fewer signals dominate, the IMD3 energy approaches the two-tone response.

For the Chicago collection, the standard deviation of the IMD_3 estimate is 5.8 dB. This corresponds to an input energy standard deviation of 1.78 dB. The energy detected on any single frequency has a standard deviation of 1.49 dB, so the error in estimating the intermodulation energy is only slightly more than the inherent uncertainty in knowing the exact energy distribution in the channel. This lends confidence to the computation of noise energy. The remaining step is to determine what portion of the energy falls within the signal bandwidth (b_0). This energy is distributed primarily across the range of the filter bandwidth ($BW\,f_c$). This calculation is practical for even low performance processors and energy limited platforms.

In summary, the effect of front-end overload is significant for many of the environments in which a cognitive (or non-cognitive) radio will operate. It is possible to establish a straightforward and readily computable polynomial relationship between a low-resolution measurement of total energy in each of the front-end filter pass bands and a high confidence estimate of the total energy that will be distributed across the signal bandwidth of interest.

4.2 Front-end Linearity Adaptation Evaluation Metrics

Performance assessment of front-end linearity adaptation algorithms will be in terms of the environmentally induced Probability of Front-end Overload ($P_{\text{Feoverload}}$) function and Intermodulation Induced Mean Noise Floor (IIMNF). The first function quantifies the likelihood that a given algorithm and design will experience overload in any specific environment (as described in Chapter 3), and the second determines the likely noise floor induced by intermodulation levels well below the IIP3 value.

To assess the effectiveness of these techniques, the assessment criterion compares methods that create equivalent reliability and performance in the identical environments, through both cognitive and non-cognitive mechanisms.

4.2.1 Probability of Front-end Overload

The first metric is probability of overload, which combines the effect of the actual spectral power distribution, the effect of the filtering process on this distribution, and the overload characteristics of the front-end, as developed in Section 4.1. From Equations (3-1) thru (3-4), the probability that a front-end will be overloaded can be expressed as:

$$
\begin{aligned}
P_{\text{Feoverload}} \quad &= \mathcal{P}\,(P_{\text{in}} > IIP3) \\
&= 0 \text{ if } FE_{\max} < IIP3, \text{ otherwise} \\
&= 1 \text{ if } FE_{\min} > IIP3, \text{ (unlikely) otherwise} \\
&= 1 - I_{x\text{IIP3}}\,(FE_{\alpha}, FE_{\beta}), \quad \text{where:} \\
&\quad x_{IIP3} = \frac{IIP3 - FE_{\min}}{FE_{\max} - FE_{\min}} \\
&\quad \text{FE}_{\min} \leq \text{IIP3} \leq \text{FE}_{\max}
\end{aligned}
\tag{4-10}
$$

where:

I_{xIIP3} is the regularized incomplete Beta function

$FE_{\min}$, $FE_{\max}$, FE_{α} and FE_{β} are characteristic of the spectrum environment, for a given value of BW

P_{in} is the power into the LNA (from the pre-selector)

$P_{\text{Feoverload}}$ is the probability of Front-end Overload

The overload probability distribution for a non-cognitive radio is shown in Figure 4-5 for a range of IIP3 and bandwidth values. This function is clearly consistent with operational experience, and with notional expectations of the relative importance of IIP3 performance and filter bandwidth in overload performance. Although the IIP3 level is used as the value of the overload threshold, the threshold can be established at any input signal level that would result in significant receiver compromise, and loss of utility.

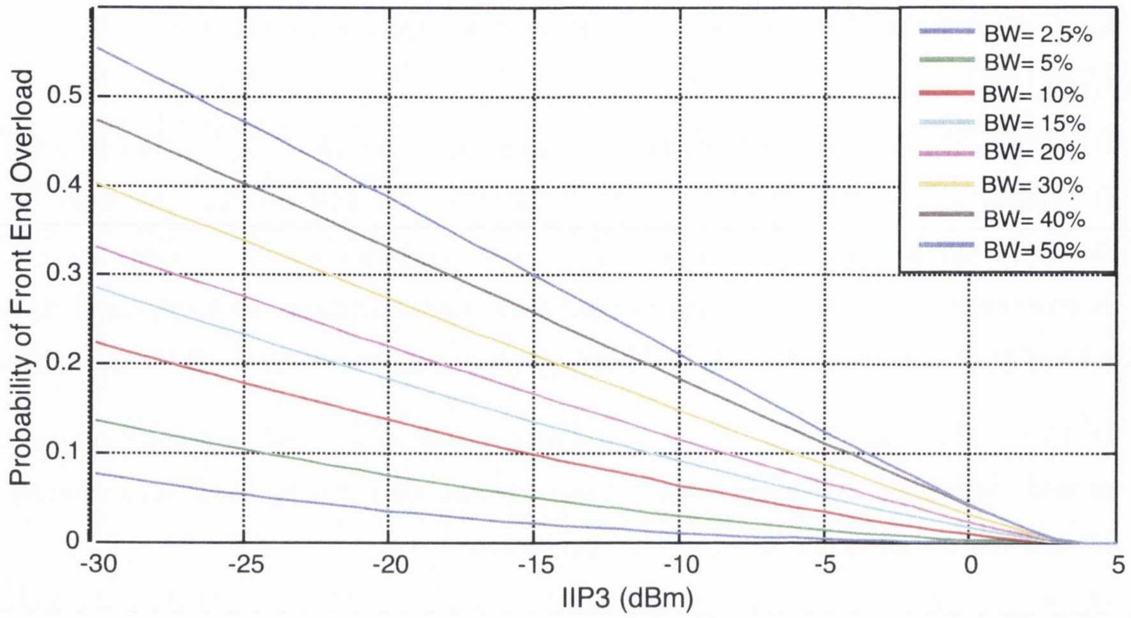

Figure 4-5 ***Non-Cognitive Radio $P_{FEOverload}$ for a Range of IIP3 and Filter Bandwidths***

In this model, only the third order distortion is considered. However, if the bandwidth of the signals provided to the LNA was sufficiently wide (more than one octave), then the function would have to reflect the IIP2 performance. However, an octave filter width in a cognitive radio would be a poor performer in most of the environment investigated (all but the NRAO). Although these results appear severe, they match anecdotal experience with wideband receivers in the presence of strong broadcast signals, such as TV or FM radio, or in co-site situations with lower energy emitters.

The non-cognitive radio baseline considers the case in which the radio is assigned to any frequency within an operating range, and the filter is adjusted accordingly. The percentage of time the power in the filter distribution exceeds the limit of the LNA's linear limits (without desensitizing it by AGC, which would reduce adjacent channel sensitivity) is $P_{\text{Feoverload}}$. Since overload conditions are generally long compared to symbol time, this measure should be considered in the context of the Packet Error Rate (PER) of the link, or the probability of even closing the link, and therefore is a reliability metric. It would not be appropriate to include this probability within a BER budget, due to the temporal correlation.

It is clear that even high levels of LNA performance (specifically IIP3) are not adequate to assure high reliability and performance communications in bands that do not have homogenous usage, such as cellular up/down links, satellite links, or other spectrum that has been segregated for "likes with likes" [136]. Section 4.3 will compare cognitive radio algorithms against these performance benchmarks and determine the required IIP3 to support identical $P_{Feoverload}$ values.

4.2.2 Intermodulation Induced Front-end Noise Elevation

Even if the front-end is not driven into a completely distorting region, the intermodulation noise can be a significant contributor to the noise floor encountered by the receiver. From (4-9), for a given bandwidth and IIP3, the total overload energy is the composite function of the distortion energy function, and the energy distribution function. The intermodulation noise induced by various performance front-ends is shown in (4-11).

$$P_{IMD3} = \frac{b_0}{BWf_c} k_1 P_{in} - 2IIP3 + k_2 \quad \text{(in dBm)} \tag{4-11}$$

Solving for P_{in} provides:

$$P_{in} = \left(\frac{BWf_c}{k_1 b_0}\right)\left(P_{IMD3} + 2IIP3 - k_2\right) \tag{4-12}$$

Substituting this into (4-10), the probability distribution of the IMD3 product is given by solving the Beta distribution:

$\mathcal{P}(\text{IMD3}_{NCR}) = I_{xIMD3}(FE_\alpha, FE_\beta)$ for $0 \le x_{IMD3} \le 1$, where:

$$X_{IMD3} = \frac{\left(\frac{BWfc}{k_1 b_0}\right)\left(P_{IMD3} + 2IIP3 + k_2\right) - FE_{\min}}{FE_{\max} - FE_{\min}} \tag{4-13}$$

and

FE_α, FE_β, $FE_{\min}$ and $FE_{\max}$ are characteristic of the spectrum environment, and BW value

$IMD3_{NCR}$ is the IMD3 noise in a non-cognitive radio

The third order intermodulation energy for the specific case of a -5 dBm front-end in the Chicago environment is shown in Figure 4-6 for a 25 kHz signaling bandwidth[1]. The dotted line represents the point at which the noise energy generated in the front-end exceeds the level of the typical receiver background noise (≈ -170 dBm/Hz), and the link margin would begin to be impacted. For purposes of demonstration, the figure depicts operation centered at 1.5 GHz, and with an IIP3 value typical of a quality consumer-level device. These values can be readily adjusted to other IIP3s by translating the curves to account for the difference in IIP3.

---

[1] This absolute signaling bandwidth is provided for comparison to typical narrow band signaling applications, such as FM voice. The analysis could have been conducted in units of bits/Hertz, but this would not reflect the non-independent nature of correlations in bandwidth usage.

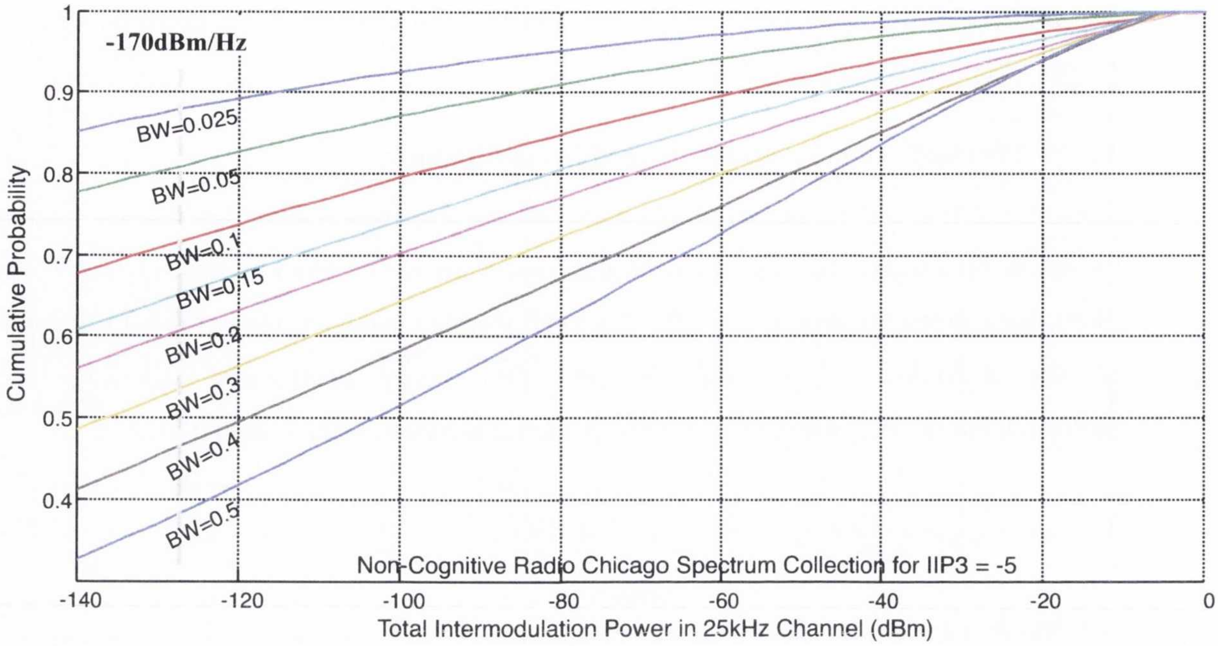

Figure 4-6 ***Non-Cognitive Radio Front-end Intermodulation Induced Noise Floor Elevation for IIP3 = -5 dBm in Chicago Spectrum***

The figure can be read as providing the probability of less than the horizontal axis level of energy for a given front-end performance level. Traditionally, link planners allocate margin for potential in-band interference, and other link degrading phenomena, but they rarely allocate margin to address the internally generated artifacts from channel intermodulation within the radios front-end. Yet, even with reasonable performing filters that have a 10% bandwidth, there is a 20% chance of 10 dB increase in noise floor if randomly assigned throughout the spectrum.

A reasonable reaction to this figure might be that radios do not typically fail as often as this metric would imply. In fact, the effect of front-end noise is often not recognized, or addressed by relocation of nodes to resolve "co-site" (Note that a cognitive radio cannot request other non-cooperating radios to physically move!). Another reason that it is less common in daily use is that in the non-cognitive radio environment, spectrum allocations and service architectures inherently avoid this situation. For example, cellular handsets only have to receive the closest base station, and do not have to listen on the same frequencies that the surrounding handsets are transmitting on. The up and down link frequencies are sufficiently separated so that fixed filters can reject the energy from nearby handsets that would otherwise overload the relatively low dynamic range front-ends. However, when contemplating ad-hoc, peer to peer, and cognitive radios, these architectural protections are not present. A node needs to be able to hear the nodes distant from it, and can reasonably expect to be immersed in other nodes and networks doing the same thing. The transition to a more ad-hoc and opportunistic vision of wireless networking will force us to abandon the architectural protections of today's services and rely on either greatly increased component performance or adaptation technologies.

4.3 Front-end Linearity Management Algorithms and Methods

The algorithms for front-end linearity are in a way an extension of the ones for dynamic spectrum. While the dynamic spectrum algorithms look for available bandwidth within the bandwidth of the pre-selector, the front-end linearity management function must look for pre-selector bands that have the minimum energy within the tuning range of the receiver. In the case of front-end linearity, the bands are defined by the performance characteristics of the pre-selector filters. This section is based on the assumption that the linearity of the receiver is driven by the wideband RF stages prior to the first mixer stage, or the A to D conversion stage, since these are the primary drivers of end-to-end IIP3.

The radio in a non-cognitive mode is assumed to operate on uniformly distributed over the pre-selector center frequency choices. The cognitive radio is assumed to be able to select the most apparent optimal operating point, and therefore, the number of statistically independent choices are dictated by the filter resolution (essentially its bandwidth) and the tuning range of the device. For example, if the filter bandwidth was 10%, and the radio tuned over two octaves, then there would be approximately 15 discrete and statistically independent choices[1].

4.3.1 Pick Quietest Band First Algorithm

The baseline spectrum density mitigation algorithm is "*Pick Quietest Band First*" or PQBF. Intuitively, the algorithm consists of applying all possible filter selection choices to the incoming (unfiltered) spectrum and measuring the total power. We imagine that the cognitive radio has a set of filter tuning parameters, with discrete center frequencies and relatively constant effective Q factor. This is certainly a reasonable assumption for typical varactor or Microelectromechanical Systems (MEMS) [137]-[138][139]. Sole reliance on high and low pass filters, or fixed band-pass, is not considered, as their performance is generally unacceptable in any sophisticated conventional or cognitive radio, and they provide no options for adaptation.

---

[1] The independence of this selection is somewhat constrained by the statistical characteristics of the spectrum assignment process. For example, if one pre-selector band is tuned to a TV broadcast signal, it certainly raises the probability that the adjacent one is also a TV band.

PQBF simply locates the lowest energy filter option and therefore, the minimum front-end linearity challenge. By implication, this decision is reached on the instantaneous (or very short integration time) energy of the spectrum that was sensed. The number of filter settings (in an idealized, fully accessible, and de-correlated environment) is given by:

$$\text{Pre-Selector Settings} = PSS = \frac{\log\left(\dfrac{f_{high}}{f_{low}}\right)}{\log\left(\dfrac{1+\dfrac{BW}{2}}{1-\dfrac{BW}{2}}\right)} \tag{4-14}$$

In fact, the number of pre-selector settings that are likely to be practically, or effectively, available are typically less than this due to limitations on the use of spectrum, and correlation of usage (some pre-selector tuning bands may be less likely to be usable if their neighbor is unusable). This can be reflected in reducing the value of effective *PSS* accordingly, to reflect local and frequency usage specific conditions.

The probability that all of the filter settings contain a signal that is above the *IIP3* threshold is the overload probability of a non-cognitive radio (which is the overload probability of a single pres-selector setting) to the power of the number of settings, given by[1]:

$$P_{FeoverloadCR} = P_{FEoverloadNCR}{}^{PSS} = (1-I_{xIIP3}(FE_{\alpha}, FE_{\beta}))^{PSS}$$

Where:

$$x_{IIP3} = \frac{IIP3 - FE_{\min}}{FE_{\max} - FE_{\min}}, \quad \text{and } FE_{\min} \ll \text{IIP3} \ll FE_{\max} \tag{4-15}$$

$P_{\text{FEoverloadNCR}}$ is the probability of front-end overload for $PSS = 1$

The performance of the PQBF algorithm is shown in Figure 4-7 for the same spectrum data sets and filter bandwidth values as in Figure 4-5 for one octave of tunable range[2]. Other tuning ranges can be reflected, since the *PSS* metric is unit-less. Note that this allows significant reduction in the level of IIP3 required for equivalent levels of overload probability. The effect of filter bandwidth is compounded, since the filter bandwidth has a significant effect on the amount of energy admitted to the front-end, and as the filter

---

[1] The CR and NC suffixes denote Cognitive Radio and non-Cognitive Radio respectively.

[2] The performance of cognitive radios with reasonable filters is convergent on essentially zero intermodulation effects after 1 octave of tuning range. For poorly performing filters, the results can be extended through interpolation of results, using Equation (4-14) to determine the equivalent value of PSS. A one octave tuning range will be used throughout the examples in this chapter.

becomes narrower, it also provides more trials from which to select a band without excess energy. Note that several of the bandwidths do not appear within the range of the graph because the values of $P_{Feoverload\mathrm{CR}}$ are below the lowest axis (10^{-10})!

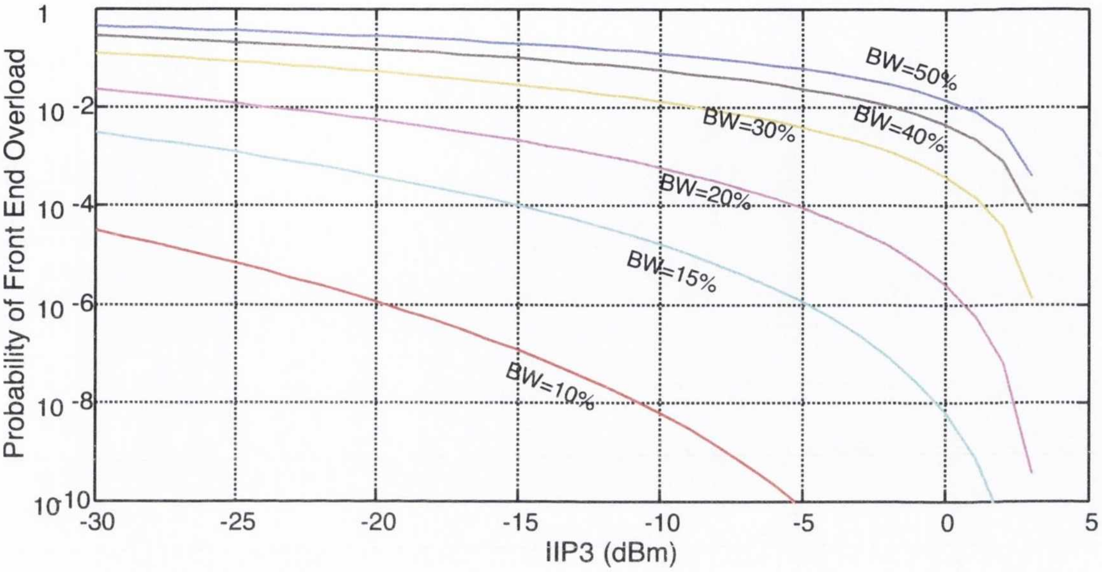

Figure 4-7 ***Probability of Overload for Cognitive Radio using Pick Quietest Band First Algorithm Over One Octave***

Inspection of these results leads directly to the conclusion that the most important resource to avoid front-end overload in a cognitive radio is the filter: increases in IIP3 performance are much less significant in reducing the probability that a given node will be overloaded. There is no obvious comparison between a cognitive and non-cognitive radio; the exponentially better performance of the cognitive radio cannot be achieved by any reasonable non-cognitive alternative. For example, for an IIP3 of -5 dBm and a filter bandwidth of 20%, the non-cognitive radio has a four percent probability of overload, while the equivalent cognitive radio would have one of 10^{-7}.

A similar probability distribution can be developed for the distribution of intermodulation induced noise floor elevation. In this case, the probability is the possibility that for a given noise threshold level, none of the filter selections had lower energy. This is the binomial probability that all choices were above a given threshold. This is Equation (4-13), raised to the power of the number of pre-selector settings.

$$\mathcal{P}\,(P_{\mathrm{IMD3}}) = (I_{x\mathrm{IMD3}}\,(FE_{\alpha}, FE_{\beta}))^{PSS} \text{ for } 0 \le x \le 1, \tag{4-16}$$
where all symbols are as previously defined

Figure 4-8 illustrates the distribution of noise floor elevations for a set of representative component performance levels in a 25 kHz signaling bandwidth. The dotted line represents an equivalent noise floor of -170 dBm/Hz.

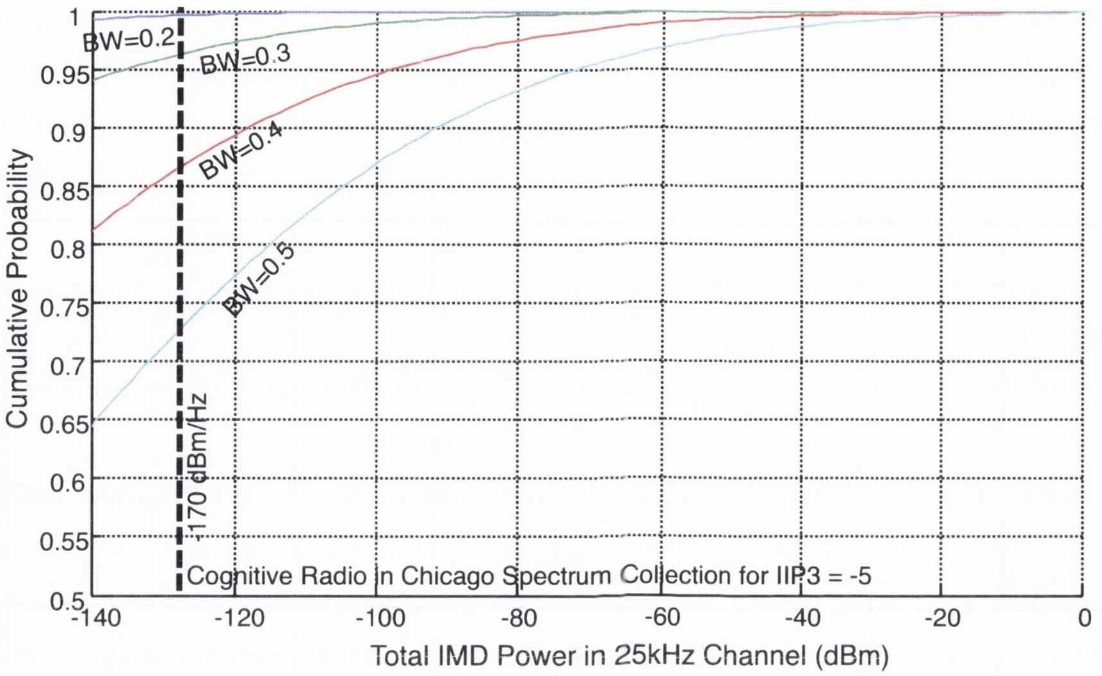

Figure 4-8 Probability Distribution of Intermodulation Induced Noise Floor Elevation when using Pick Quietest Band First Algorithm

The cognitive radio has a high probability of successfully locating spectrum that has minimal intermodulation noise, given that it has both a reasonable filter selectivity and sufficient frequency range to provide pre-selector candidates. This performance is essentially unchanged for levels of IIP3 as low as -20 dBm. Filter bandwidths better than 30% show low IIP3 product noise even with poor IIP3 performance, even in the densest environment investigated.

Radio designers have generally been reluctant to include high selectivity filters in front of the LNA stages in order to avoid the noise figure degradation that is a consequence of the filter insertion loss. It is clear that in dense and stressing environments, the use of high selectivity filters has significant benefits for both reliability of operation and for noise floor reduction, despite the signal attenuation that this design choice might cause. It is more predictable to allow for a low level of fixed insertion loss, than to provide adequate margin for the range of potential intermodulation noise floor elevation.

4.3.2 Marginal Noise Impact Consideration

A more complex algorithm recognizes that while there is certainly a unique lowest energy band, there are probably a number of selections that are essentially indistinguishable from each other in terms of intermodulation product energy elevation of the noise floor, since, if the IMD3 noise is below the natural noise floor, its impact is minimal. The Consider Marginal Noise Impacts (CMNI) algorithm defers selection of specific bands, and provides the upper network layers with a set of alternative choices, along with the likely noise floor impact of each choice. Figure 4-9 shows the probability of various levels of noise for the Chicago sample, for an IIP3 of -5 dBm, and an octave of coverage.

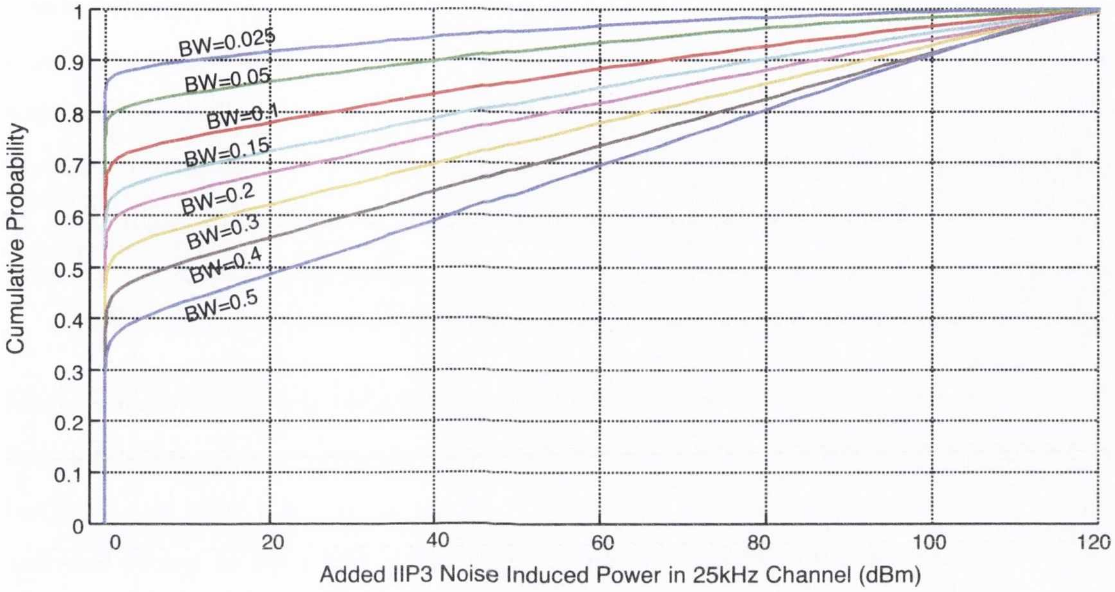

Figure 4-9 Probability of Additional Intermodulation Induced Noise

Even the worst filter has a range of choices that are within a single dB of the noise floor, and a further range of 5% or so with slightly elevated noise levels, but which still could be attractive for other reasons (antenna performance, propagation, etc.). The overload and noise impact performance of this algorithm is essentially unchanged from the PQBF algorithm. It has the advantage that it can utilize the mapping of input energy to noise energy described previously, and characterizes a set of choices that a network could make without significant processing cost or inaccuracy.

Substituting an acceptable noise floor impact for $IMD3_{CR}$ in (4-16) yields the probability that the noise floor of any given pre-selector will be below the threshold, and is therefore the ratio of pre-selector choices that qualify for use. For example, setting this equation equal to the noise floor would result in the ratio of choices yielding a maximum 3 dB increase in the noise floor for the cognitive radio. In contrast to PQBF, this algorithm fully implements the philosophy and structure of Section 2.4, in that it defers the choice of frequency until the considerations from other layers are concatenated, and an overall performance metric can be determined. As such, it serves as an extensible cognitive radio building block that can address upper layer considerations of range, power and energy, topology, and reliability in a more globally optimized network structure.

4.4 Front-end Linearity Management Benefits

There are two measures of benefit for front-end linearity that will be examined. The first is the reduction in noise floor that can be achieved by selecting filter settings that minimize the noise generated in the receiver's front-end, and the second is the corresponding ability to reduce the performance of the receiver's components, and still maintain equivalent performance. These are two aspects of the same functionality. A high performance application might look to cognitive radio technology to resolve issues of reliability that are induced by co-site interference, while more cost sensitive applications might look to cognitive radio to enable lower cost or energy consumption devices to perform equivalently to non-cognitive and non-adaptive designs.

The performance effects of linearity management are highly dependent on the signal environment. Underneath a cell tower, the downlink may have a very uniform signal density. On the other hand, in a Public Safety band, there may be a great mix of signal strengths. In the following discussion, an initial baseline was used to assess benefits, which were centered on the performance region depicted in Table 4-1.

Table 4-1 Notional Front-end Analysis Performance Points

Measure	Symbol	Value	Units	Basis	Cognitive Radio Treatment
Filter Bandwidth	*BW*	20	%	Moderate performance tunable filter	Not varied in analysis of CR performance
Pre-IF Gain (Power)	G_{LNA}	10	dB	Low performance LNA	Not varied in analysis of CR performance
IIP2	*IIP2*	40	dBm	>> Input Energy	Not considered due to filter coverage
IIP3	*IIP3*	-5	dBm	Typical low cost LNA	CR performance reduced until equivalent performance results

The frequency of operation of the radio in a non-cognitive mode is assumed to be static, and assigned to random channels that are evenly distributed over the operating range of the device. The cognitive radio is assumed to be able to select the most optimal operating point. This relationship will be used to determine the energy probability distribution to which the front-end will be subjected to for any given spectrum power distribution.

4.4.1 Probability of Front-end Overload

The detrimental effect of front-end overload is significant. It is highly likely that an appropriately sized communications link will either fail to acquire, and/or fail to achieve operation. To understand the performance benefits offered by cognitive radio, the improvement offered by the adaptations can be expressed as:

$$\mathrm{Benefit}_{\mathrm{Poverload}} = \frac{P_{\mathrm{FEoverloadNCR}}}{P_{\mathrm{FEoverloadCR}}} = \frac{1 - I_{x_{\mathrm{IIP3}}}(FE_{\alpha}, FE_{\beta})}{(1 - I_{x_{\mathrm{IIP3}}}(FE_{\alpha}, FE_{\beta}))^{PSS}} \tag{4-17}$$

which simply reduces to:

$$Benefit_{\mathrm{Poverload}} = (P_{FE\mathrm{overloadNCR}})^{-PSS+1} \tag{4-18}$$

These performance relationships are quite extreme. Some illustrative curves of the benefits of adaptation are shown in ***Figure 4-10.***

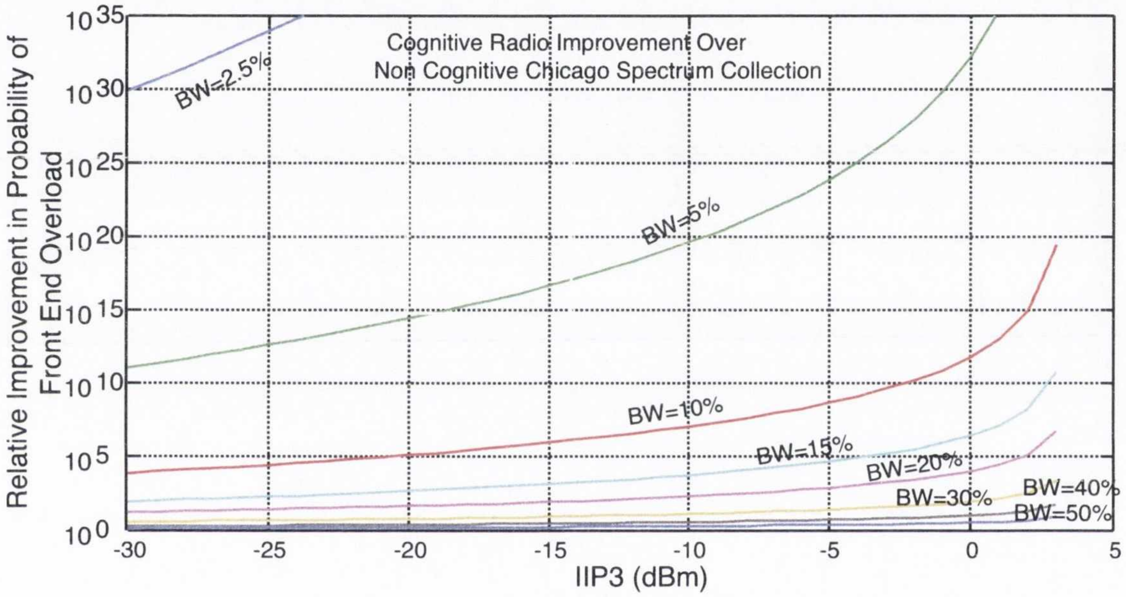

Figure 4-10 Illustrative Improvement in Probability of Front-end Overload

For the one filter setting option, (*PSS* = 1) the performance is identical. The benefits of the adaptation mechanism are directly related to the performance of the non-adaptive radio. If the chance of overload in a given environment was 5%, then with ten independent filter choices (*PSS* = 10) the chance of overload in this environment is only $(0.05)^{10}$, for a benefit improvement of $5\mathrm{x}10^{11}$. The values of *PSS* can be adjusted to reflect regulatory or other constraints on the availability of spectrum.

Another modality to exploit adaptation is to examine the reduction in *IIP3* that provides equivalent performance using adaptation. This is an equally important consideration as performance, since it maintains equivalent performance levels, even with affordability due to reductions in the component performance levels[1]. To analyze the component performance reductions, the performance of both approaches are set equal.

$$P_{\text{FeoverloadNC}} = P_{\text{FeoverloadCR}} = (I_{x\text{IIP3CR}}\,(Fe\alpha,\ Fe\beta)) = (I_{x\text{IIP3NCR}}\,(Fe\alpha,\ Fe\beta))^{PSS} \qquad (4\text{-}19)$$

Where:

x_{IIP3CR} and x_{IIP3CR} reflect the IIP3 values of the cognitive and non-cognitive radios respectively

A symbolic solution to this equation is difficult, but computation of $P_{\text{Feoverload}}$ for a range of IIP3 values and pre-selector settings provides a reasonable understanding of how this adaptation lowers the IIP3 requirement for a cognitive radio. This is shown in Figure 4-11.

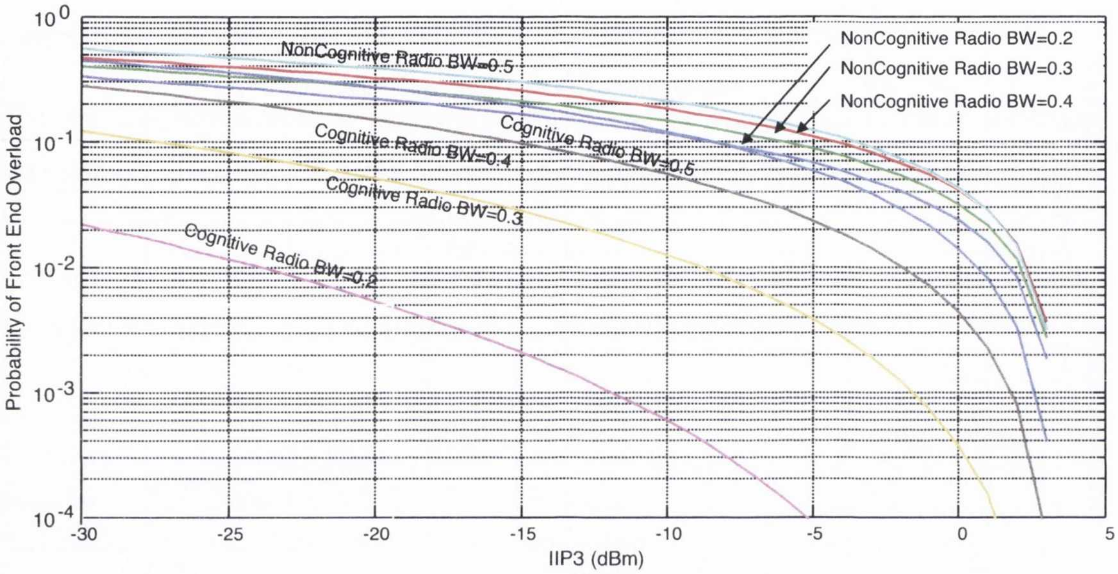

Figure 4-11 ***Values of IIP3 and PSS for Equivalent Values of $P_{FEOverload}$ in Conventional and Cognitive Radios***

A cognitive radio with only a 30% filter and IIP3 of -27 dB has the equivalent probability of overload as a non-cognitive one with -5 dBm IIP3! The performance of a cognitive radio with a 20% filter shows over 30 dB of improvement. Table 4-2 shows the metrics from the implementation of cognitive radio techniques for the component capabilities shown in Table 4-1.

Table 4-2 ***Probability of Front-end Overload***

Improvement in Performance	> Reduction from 4% failure rate to essentially zero failures
Reduction in Required Hardware Capability	> 20 dB reduction in front-end IIP3 with reduced probability of overload over most of the typical design range

---

[1] Note that trade space in reduction in filter performance is not included since it has been shown that even a non-cognitive radio with arbitrarily narrow filters can not create reliable noise floors that are the equivalent to the cognitive radio.

It would be difficult to make a cognitive radio with performance so deficient that it cannot locate at least one frequency of operation that is not subject to overload, given reasonable filter bandwidth and operating range. Although it is the intent of this dissertation to develop equivalently performing designs, it would appear that no non-cognitive radio can offer the overload avoidance performance of a cognitive radio with comparable performance components, unless the IIP3 threshold is increased well beyond the conceivable energy within pre-selector bands. In the next section, even this unbounded approach will be shown to be inadequate to avoid noise floor elevation.

4.4.2 Noise Floor Elevation Impacts of Front-end Linearity

Equation (4-11) established the performance of a cognitive and non-cognitive radio within given spectrum environments. The process for computing this performance metric is similar in approach to the one described for front-end overload, except focused on the linear effects, rather than the discrete overloaded versus not overloaded conditions examined previously. Since significantly elevated noise levels are not (or at least should not be) typical in any radio, the value of the mean or median noise elevation is not particularly insightful to understanding the benefits of adaptation. The crucial objective is to achieve equivalent confidence in performance. Therefore, the critical need is to determine the noise (and decrease in noise) for a given probability of occurrence. Solving equations (4-11), (4-13) and (4-16) yields:

$$Benefit_{FENoise}(p) = \frac{Prob_{NC}}{Prob_{CR}} = \frac{Ix_{NC}(FE\alpha, FE\beta)}{(Ix_{CR}(FE\alpha, FE\beta))^{PSS}},$$

Where:

p is a fixed probability of overload

$$x_{NC} = \left(\frac{BWf_c}{k_1 b_0}\right)\left(\frac{P_{IMD3-NC} - 2IIP3_{NC} + k2 - FE_{\min}}{FE_{\max} - FE_{\min}}\right) \tag{4-20}$$

$$x_{CR} = \left(\frac{BWf_c}{k_1 b_0}\right)\left(\frac{P_{IMD3-CR} - 2IIP3_{CR} + k2 - FE_{\min}}{FE_{\max} - FE_{\min}}\right)$$

An example of cognitive radio IMD$_3$-induced noise reduction is shown in Figure 4-12. Much of this benefit is reduction below the effective noise floor, and therefore yields no significant performance advantage. In the following analysis, a 90% performance point will be examined. This is the point where 90% of the performance is better than stated, and 10% is less. Some applications might desire a more stressing analysis basis.

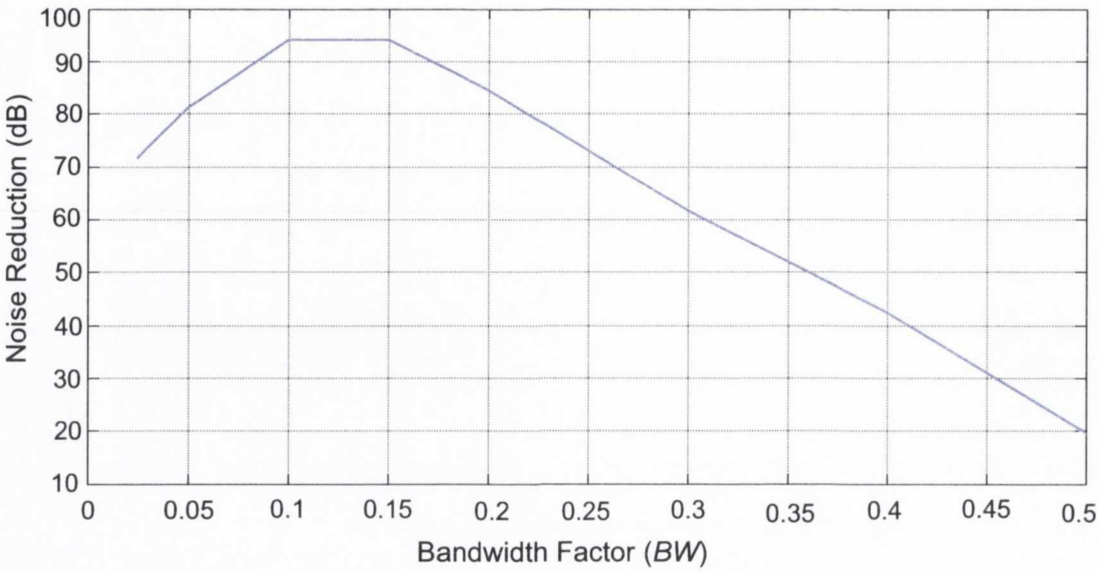

Figure 4-12 Median IMD3 Noise Reduction through Cognitive Radio

Figure 4-13 illustrates the 90% confidence level, with an assumed noise floor at -170 dBm/Hz.

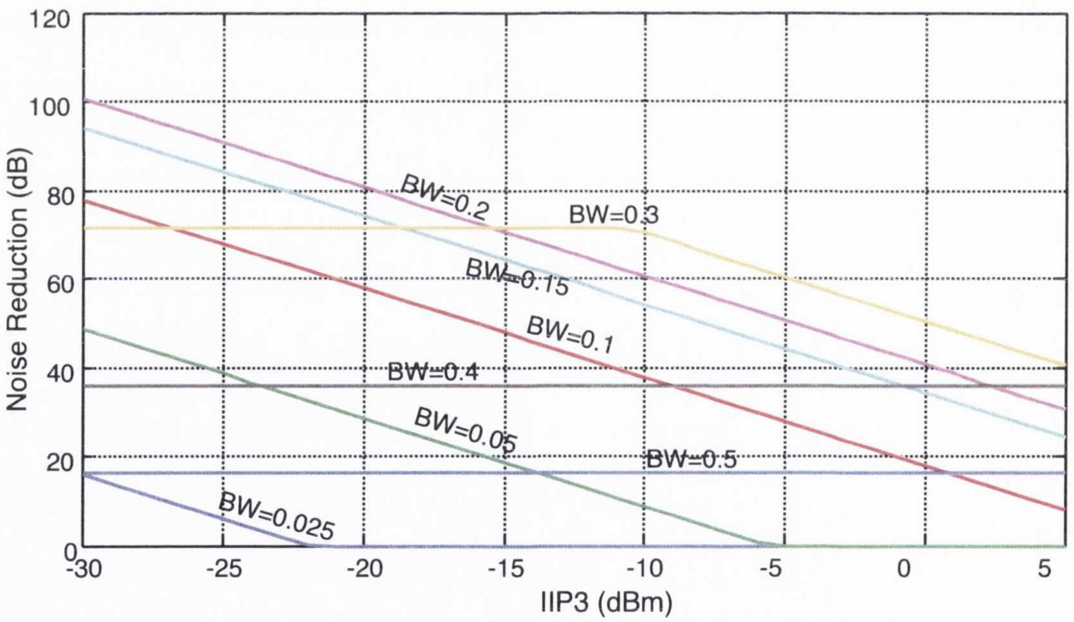

Figure 4-13 Beneficial Region of Cognitive Radio Noise Reduction in dB (at 90% performance point)

The noise energy is sufficiently high that all of the IIP3 performance levels generate IMD3 above the -170 dBm/Hz level for the 40% and 50% bandwidth filter at least 10% of the time. The 30% filter has constant benefit until the IIP3 is sufficiently high to reduce the IMD below the noise, and the extremely narrow filters (2.5% to 20%) have diminished improvement as the IIP3 is improved. The benefits of the cognitive radio adaptation become even more extreme in high performance environments. Throughout the typical operating region of -10 to -5 dBm IIP3 and with 25% filters, the benefits in the 90% case exceeds 40 dB.

The introduction of cognitive radio technology enables significant reductions in the component performance needed to meet identical performance levels. The primary interest is reduction of required LNA linearity, as reflected in IIP3 levels and in filter resolution, or *BW* factor. Solving (4-19) for identical noise ratios (*Benefit* = 1, or 0 dB) and varying IIP3 shows the anticipated reduction in component performance for LNAs as a function of the difference in the energy distribution. A numerical analysis is provided in Figure 4-14.

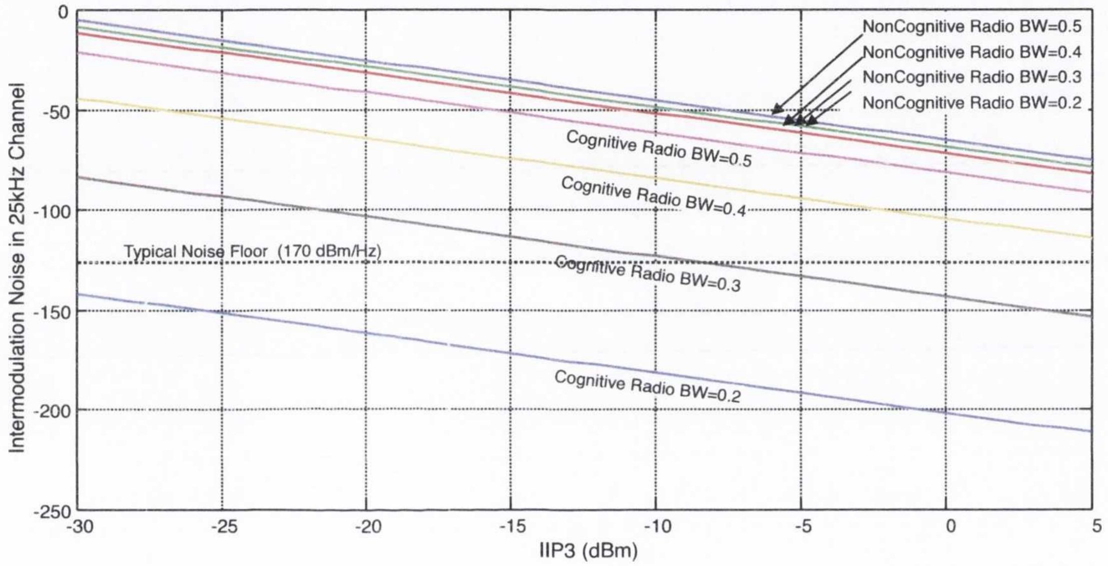

Figure 4-14 90% High Energy Equivalently Performing Combinations of IIP3 and PSS for Cognitive and Non-Cognitive Radios

This figure can be read by comparing the cognitive and non-cognitive radio filter bandwidth (*BW*) lines for identical IMD3 noise. The noise level at the 90% point for a non-cognitive radio is not highly sensitive to the filter parameter, but the cognitive radio is highly sensitive to the filter technology. The benefits of adaptation increase rapidly with even low performance pre-selector filters, or more specifically with values of *PSS* (pre-selector option which can be obtained through filter performance or octave coverage). Use of adaptation with even 30% filters yields equivalent performance at 30 dB less IIP3 performance. The sensitivity to *PSS* is clear, since the 50% filter (*PSS* = 2) only has a 5 dB advantage!

Table 4-3 illustrates reasonable expectations for anticipated impacts from the introduction of cognitive adaptation in reducing the intermodulation noise floor. These are nominal values of an IIP3 of -5 dBm and 20% filter bandwidth.

Table 4-3 Front-end Linearity Effect on Noise Floor

Improvement in Performance	The improvement is highly dependent on IIP3 and filter bandwidths, but is typically 20 to 55 dB throughout a reasonable operating range at a 10% probability of occurrence.
Reduction in Required Hardware Capability	> 25 dB decrease in IIP3 can maintain IMD3 below the typical channel noise floor with filters above 30% bandwidth.

Consideration of generalized spectrum environments is also instructive. Figure 4-15 illustrates the 90% performance benefit (in dB reduction in the noise floor) for a 10% filter bandwidth, using the monotonic spectrum indices introduced in Chapter 3.

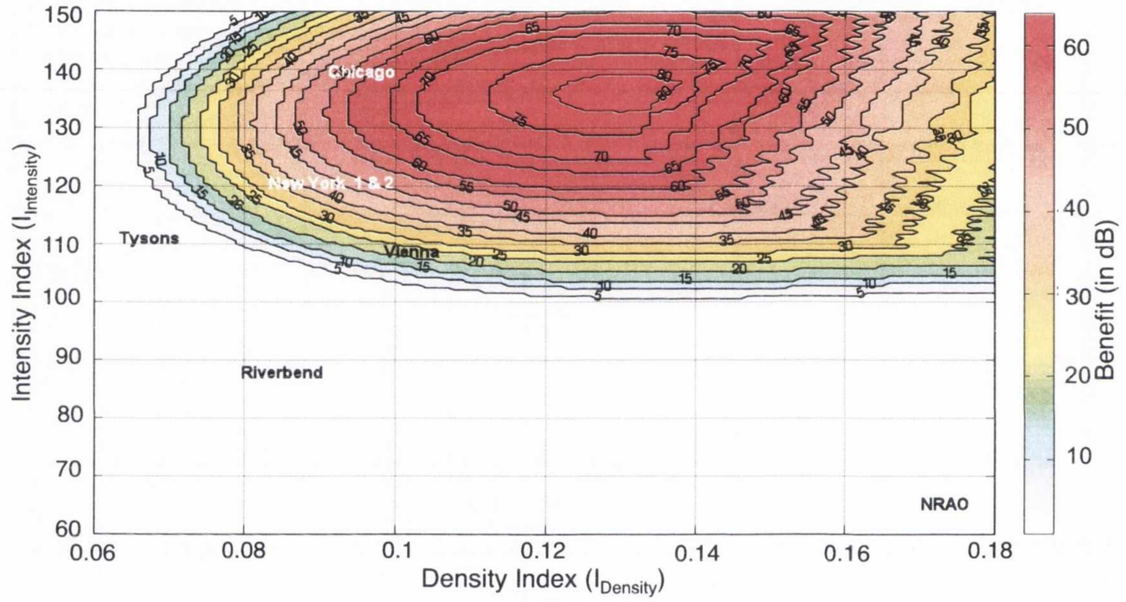

Figure 4-15 Cognitive Radio Benefits (in dB) for 20% Filter and a -10 dBm IIP3 front-end LNA in a Range of Spectrum at the 90% performance point

The benefits of cognitive radio technology are most significant where the spectrum is both energetic and dense. However, with sufficiently dense spectrum, even the adaptation of the cognitive radio has reduced benefits. The Chicago and New York spectrum actually represent environments that optimally exploit cognitive radio benefits at the operating points depicted in this chapter. As more systems enter the spectrum, spectrum will become both more dense and more energetic, moving from the current environments, towards the stressing upper right of the figure; making cognitive adaptation even more critical to future wireless systems.

As the required reliability is increased, the benefits from cognitive adaptation increase as well. Figure 4-16 illustrates the same set of cases and component performance, except for a reliability performance level of 99%. The region over which the noise level is reduced is significantly extended to lower intensity and density indices, and the degree of improvement has generally increased.

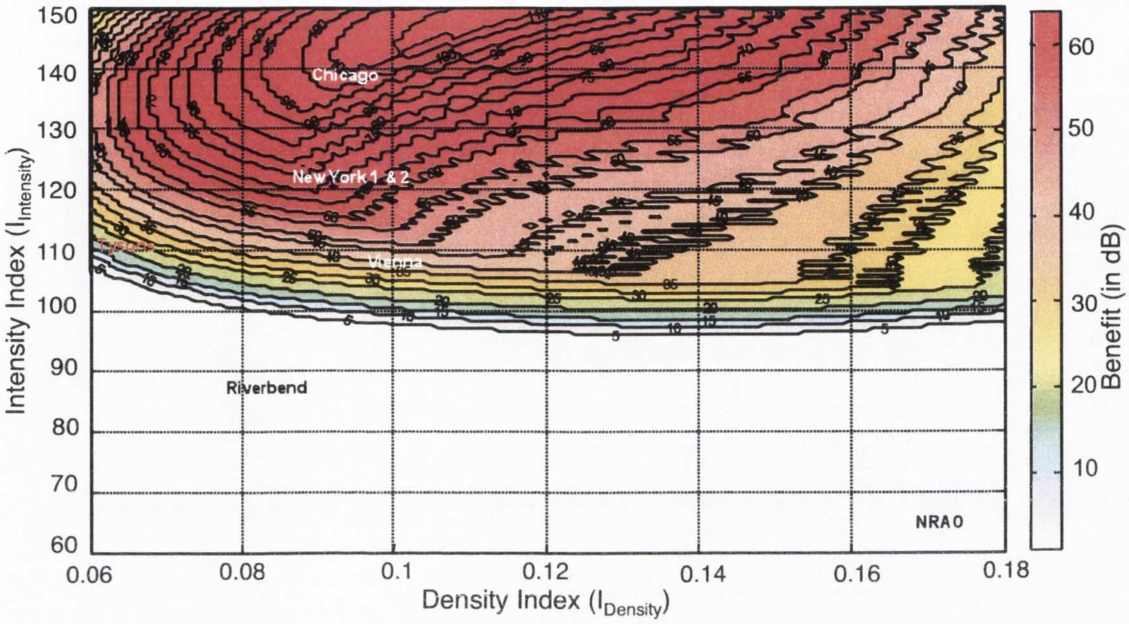

Figure 4-16 ***Cognitive Radio Benefits (in dB) for 20% Filter and a -10 dBm IIP3 front-end LNA in a Range of Spectrum at the 99% performance point.***

Even 99% is not a very stressing reliability level: it corresponds to over 80 hours of outage a year. Clearly, cognitive adaptation provides significant and enabling benefits to applications that are either in dense environments or those that must guarantee a reasonable level of service reliability.

4.5 Front-end Energy Management Conclusions

This dissertation provides a quantitative process that can yield reasonable estimates for the overload probability and noise distribution of cognitive and non-cognitive radios in spectrum bands where spectrum activity is independent and not significantly correlated (outside of the range of the tuning filters). Analysis of collected spectrum environments demonstrates that dynamic spectrum access, combined with LNA performance aware band selection, provides significant advantages in both reliability of operation, and likely intermodulation induced noise floor. These results can be extended to any spectrum environment, synthetic or real, through use of the spectrum distribution variables provided in Chapter 3.

This chapter examined the role of DSA as a fundamentally new approach to resolve interference caused by non-linear receiver response. This "internal" use for DSA has compelling advantages to system users and operators, even when operating in dedicated spectrum. This is a practical and immediate use of DSA that is not dependent on regulatory action, or establishment of new policy regimes. Once homogenous systems apply DSA to resolve their internal design issues, it may also provide the evidence for more widespread use of DSA in less controlled and more heterogeneous environments

Additionally, the front-end linearity benefits set the stage for an argument that cognitive radio is potentially a cost savings technology that can increase both performance and affordability of wireless devices, and simultaneously enable the devices to coexist in the resulting dense, highly stressed environment.

The DSA community should advance arguments for DSA that go beyond increased access to spectrum. DSA provides an effective solution to co-site and adjacent channel interference that cannot be accomplished through practical levels of linearity in receiver front-ends, and should be able to provide this mitigation at decreased equipment acquisition and energy costs. The same regulatory processes that are conservative in embracing DSA, due to its disruptive effect on spectrum regulation, may find it attractive to resolve the adjacent channel energy management issues that are exponentially more complex than the in-channel spectrum management process currently practiced. The economic consequences of adjacent channel effects are hinted at by the Nextel/public safety issues in the US [127].

Chapter 5 Local Channel Optimization using DSA

5.1 DSA Roles and Adaptive Cognitive Radio Behaviors

The next three Chapters involve successively more global scope in the application of DSA-based adaptations. In Chapter 5, the emphasis is on unilateral and bilateral, link-level, decision making to optimize individual channels, in even the most simplistic spectrum management environments. In Chapter 6, the adaptation is extended to consider the effect on other users of the spectrum, and the options available to the node and network to optimize performance within an environment of nodes, and with recognition of a responsibility to not unacceptably impact other users. In Chapter 7, DSA techniques are further evolved to require nodes to actively minimize their aggregate environmental impact through adaptation of spectrum footprint through waveform selection. These approaches are unique from, and complementary to the bulk of the DSA research literature, which has tended to address DSA from the perspective of an asymmetric responsibility for non-interfering operation to specific users of the spectrum.

The first of these objectives (and the one discussed in this chapter) is optimizing the impact of frequency selection on the noise floor encountered by a node, and/or its link partner. An important question is to what extent the process of selecting frequencies through DSA mechanisms also inherently provides the selection of low noise channels, or whether additional cognitive radio algorithms are required to further select low noise channels within a set of permitted channels. DSA interference consideration has generally focused on the effect of DSA policies on other users of the spectrum, as in the DARPA XG Program [140] principle of "do no harm". These policies may also contribute to the selection of suitable operating frequencies. For example, if the occupied spectrum threshold is 6 dB above the noise floor, in a simplistic flat earth environment, selection of the lowest noise channel could double communications range over other permissible, but less desirable channels.

For analysis purposes, two spectrum management situations are depicted in Table 5-1. These are generally characterized as a shared commons with a "right to transmit" in row A, and shared secondary usage, with protection to the primary user; representing a "right to receive" in row B. In this second case, it is assumed that the cognitive radio is provided with a set of policies that have been determined to avoid interference to the primary user.

Table 5-1 Categories of Spectrum Operation

Title	Non-Cognitive Operation	Cognitive Operation
Shared With no Constraints on Usage (A)	All radios can use the spectrum without considerations of interference, but must accept whatever the channel they use encounters.	The radio selects among all of the candidate channels for the one with the least interference impact.
Shared, with Interference Avoiding Policies (B)	A number of frequencies are determined to be suitable based on the impact on other users. This determination would be distributed between transmitter and receiver.	The list of policy compliant frequencies would be utilized in an order based on the interference effects determined for each frequency.

The primary difference is that in case *A*, the potential noise in the channel is unbounded (at least up to $SO_{\max}$) and in case *B*, the maximum energy permitted in any policy compliant band would constrain the energy that any link partner could accept on a frequency judged to be compliant.

In case *B*, spectrum is assigned through DSA, without permitting interference to the primary user(s). The radio utilizes the policy criteria to determine the specific spectrum that is allowed for the operation of a cognitive radio. A subset of frequency constraints is determined from the list of DSA policies, such as those provided by McHenry [140] and applied to available spectrum measurements to determine a policy compliant subset of these candidate frequencies. A number of different policies are possible, but most have the characteristic that the channel signal energy present on a frequency is constrained to fall below a stated threshold. The channel energy levels of candidate frequencies at the lower energy end of the cumulative distribution is provided in Figure 5-1.

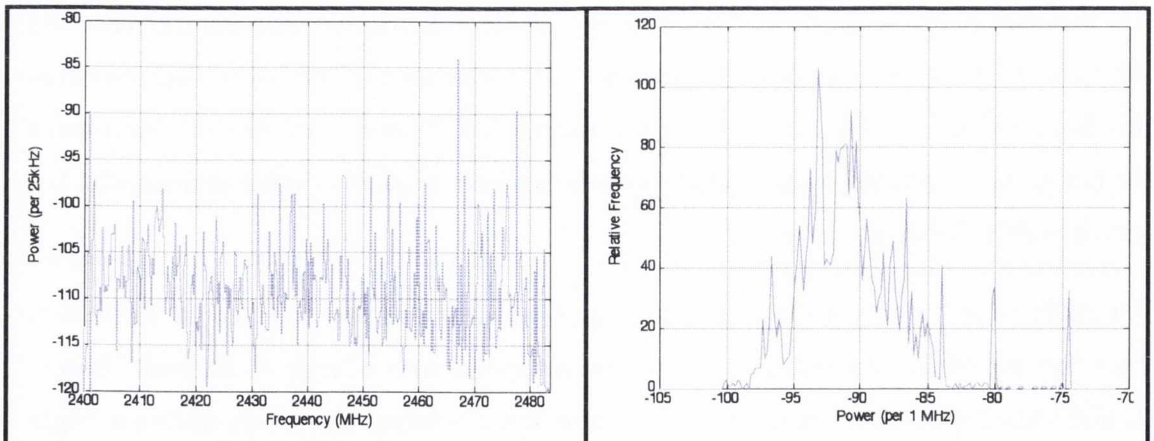

Frequency Domain

Power Histogram

Figure 5-1 Noise Distribution of Frequencies in the 2.4 GHz ISM Band in Chicago

As would be expected from the distribution established in Chapter 3, the distribution of signals is quite even in this lower energy range. An approximation of both the mean and median value is one half of the range from the noise floor to the threshold, which is typically considered to be from 6 to 10 dB above the noise. In the case of unlicensed and uncontrolled frequencies, there is not a requirement to comply with a non-interference policy, so any frequency can be used without regulatory constraint. The usage of these "common" bands is quite uneven, with some urban areas completely utilized, and rural ones are often underutilized. A typical energy distribution (Chicago, IL) for an ISM band is shown in Figure 5-2.

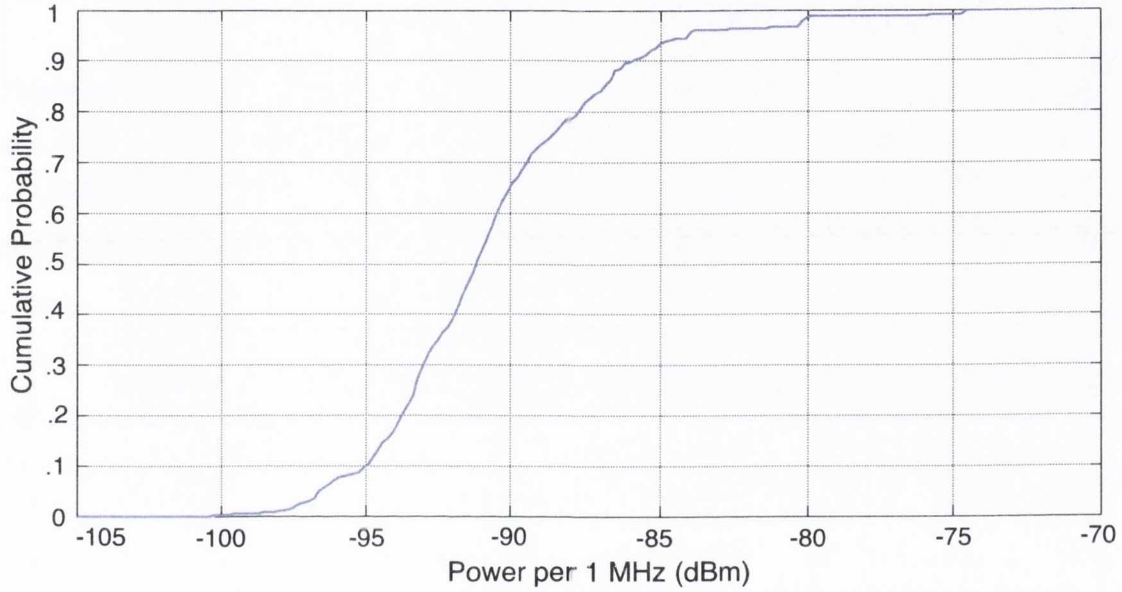

Figure 5-2 Spectrum Energy in possible 1 MHz Chicago ISM Bands (25 kHz centers)

Legally, any of these frequencies could be selected for use, since the ISM bands require no interference avoidance behavior. However, the distribution clearly shows that there are both "good" and "bad" choices available, and the "legality" of a choice is not a good criteria. Actions and decisions (for example, selecting already locally occupied frequencies) that might be harmful to other users are certainly not in the radio's own best interests. Even on a rooftop of a Chicago high-rise building, more than 50% of the choices have at least 6 dB noise floor elevation. The worst case 25% have over 10 dB noise elevation, while selection from the best case 25% would typically result in a noise floor only 2 to 3 dB above the equipment noise floor.

In policy controlled spectrum there will be a choice of spectrum that is acceptable for use, but it may have varying noise floors. Table 5-2 illustrates a notional policy set that might be enforced by a DSA device.

Table 5-2 Typical Interference Policy Controls for Constraining Spectrum Choices

Policy	Rationale
Not an ISM band (where interference policies would not apply)	Unique spectrum rules, behaviors and limits which are more constraining than other spectrum opportunities.
Signal level no more than 6 dB above state of the art sensitivity on 25 kHz bandwidths	No existing signal present, and no risk of not sensing a narrow signal using wideband integrated sensing bandwidths.
Not in restricted bands	Not a safety of life (121.5 MHz, 243 MHz), radio astronomy, or other quiet band.
Not in bands with unique directionality or transverted applications	Avoid bands where users use split frequencies, such as in satellite up/down link.

The dense spectrum environments are an appropriate model for use in understanding the "typical" environment faced by a cognitive radio. A histogram of policy compliant (in channel energy below a typical fixed threshold) spectrum in the denser spectrum collections (Chicago, New York, and Vienna) is provided Figure 5-3.

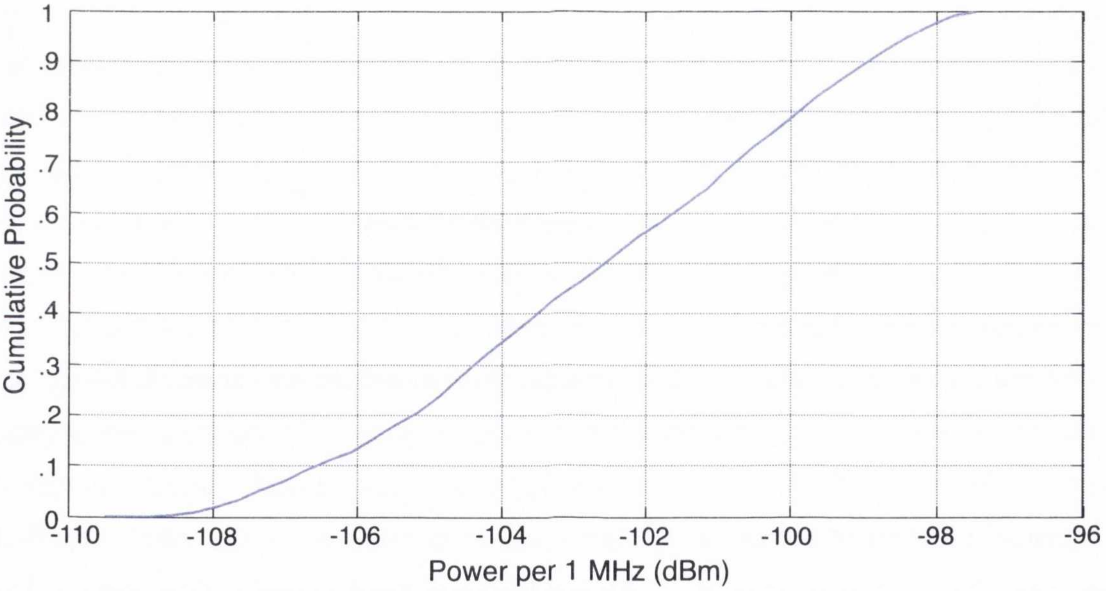

Figure 5-3 A Cumulative Distribution of Policy Compliant Spectrum in Dense Environments

The distribution is what would be expected at the lower portion of a larger Gaussian distribution. The median value of noise in this sample is -103.5 dBm. However, 10% of the available 1 MHz bands have noise floor values below approximately -107 dBm. Even within the policy compliant spectrum, there is a 4.5 dB difference in noise floor that can be exploited by a cognitive radio simply through selecting frequencies after the policy compliance process has been performed on all candidate frequencies (in contrast to accepting the first compliant candidate).

5.2 Noise Floor Reference Evaluation Metric

The metric for evaluating mean noise floor reduction is the ratio of the mean case of noise over the best choice within the permitted spectrum of operation. This is the:

$$\text{Noise Floor Impact} = \frac{\text{Median Noise Power}}{\text{Achieved Noise Power}} \tag{5-1}$$

Where:

Median Noise Power is the average noise energy in a selected signaling channel

Achieved Noise Power is the best noise level that can be achieved through selection of optimal signaling channels

This approach is one element of cognitive radio that already has been applied to a number of consumer devices, even if the technology is not directly referred to, or even recognized as cognitive radio. Cordless phone hubs scan the set of channels available to pick unoccupied channels and automatically set the handsets to the selected channel. Some Wi-Fi hubs also scan for the quietest channels before sending out SSIDs. In fact, this simple adaptation has such significant advantages, and can be performed at such low cost that it can be considered the first cognitive radio feature that has been widely deployed. Although the DSA process inherently looks for low noise levels as a byproduct of primary interference avoidance policies, most contemplated policies will permit operation even with a significant level of noise in the selected channel.

5.3 Noise Floor Management Algorithms and Methods

The Noise Floor algorithm has a simple objective and a correspondingly straightforward set of algorithm candidates. In ***Pick Lowest Noise Floor First*** (PLNFF), the radio(s) selects among candidate frequencies to select the one that has the lowest level of noise energy within a set of permitted frequencies. The radio, having performed DSA policy screening, would then further exploit the spectrum information on hand by selecting the least noisy of the policy compliant frequencies of operation, as a logical extension of the DSA process. Both are similar, although they are driven by different objectives. The DSA algorithm selects quiet frequencies in order to reduce the noise caused to other users, with an objective of "community good". The Noise Floor algorithm performs a similar process in order to achieve an effect that is intended to benefit the cognitive radio link only.

We can recognize three variant implementations of this algorithm. These are Transmitter Selects Lowest Noise Floor First (TSLNFF), Receiver Selects Lowest Noise Floor First

(RSLNFF), and Jointly Select Lowest Noise Floor (JSLNFF)[1]. In the case of duplex operation on a single frequency, the first two of these (TSLNFF and RSLNF) become equivalent, and can only be differentiated when the bilateral link frequencies can be established independently. Join selection (JSLNF) is applicable only to bi-directional (or n-directional) links, such as Wireless Local Area Network (WLAN) applications.

The functional performance differences between these algorithms directly relate to the degree of correlation of the noise floor seen by the two ends of the link. If the other users of the same spectrum are significantly distant (physically) as compared to the members of the cognitive radio link, then (in a flat earth environment) it would be expected that the signal environment would be perceived to be highly correlated. If, on the other hand, the cognitive radio has a link range that is much greater than the physical separation from other users of the frequency, then each end of the link would be in unique environments, and might have similar statistical characteristics, but would not be correlated in terms of individual frequencies.

The degree of correlation that is seen at each end of a cognitive radio link can be analyzed by randomly distributing incumbent signal sources and receivers and cognitive radio link pairs over a region. The generalized analytic approach considers three nodes: a transmitting cognitive radio, a receiving cognitive radio, and a non-cooperative transmitting node, as shown in Figure 5-4.

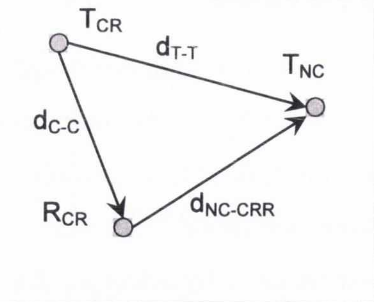

T_{CR}	Transmitting Cognitive Radio
R_{CR}	Receiving Cognitive Radio
T_{NC}	Transmitting Non-Cooperative Radio[2]
$d_{C\text{-}C}$	Distance between the two Cognitive Radios
$d_{T\text{-}T}$	Distance between the two Transmitting Radios
$d_{NC\text{-}CRR}$	Distance between the Cognitive Radio receiver and the non-cooperative transmitter

Figure 5-4 Model for Receiver Correlation as a Function of Relative Range and Propagation

It is necessary to determine if decisions made unilaterally by either endpoint are significantly less advantageous than ones made bilaterally/multilaterally between the endpoints. The signals conditions are those at T_{CR} and R_{CD} for all positions of T_{NC} for

---

[1] The relative designation of receiver or transmitter is arbitrary, but is assigned based on considering the node that first initiates link formation as the "transmitter".

[2] A non-cooperative node is any node that is not participating in the decision to utilize a frequency. It could be a non-cognitive radio node, a broadcast node, or even a cognitive radio that can not, or will not, join in the exchange of spectrum occupancy information, and which can not be assumed to have any interference avoidance behavior itself. As such, it can only be treated as a fixed and non-responsive, source of interference.

which the signal from T_{NC} at T_{CR} and R_{CR} are within the compliance threshold for a channel. If T_{NC}'s signal is above the threshold at one or both of the cognitive radio endpoints, then it would reject the frequency in any case. The basis for algorithm effectiveness is the degree of similarity of the perceptions of T_{CR} and R_{CR} regarding the interference potential of an otherwise compliant (below threshold) channel. In a non-attenuative and similarly propagating environment, the difference in signal perception is given by:

$$\left(\frac{S_{TCR}}{S_{RCR}}\right) = \frac{N_0 + K_{path}\, P_{NC}\, \left(d_{T-T}\right)^{-\alpha}}{N_0 + K_{path}\, P_{NC}\, \left(d_{NC-CRR}\right)^{-\alpha}} \qquad (5\text{-}2)$$

where :

$\alpha =$ is the propagation exponent (typically $2 \leq \alpha \leq 4$)
N_0 is the total receiver and environmental noise, assumed equal at both ends of the link
P_{NC} is the power of the non-cooperative node
K_{path} is the propagation loss for a unit distance

The constant term (N_0) provided by the detector noise floor establishes that it is of little consequence if there is great difference in the interfering signal level, if that level is well below the noise floor. To determine the link margin effects of the selection of frequencies, the noise term needs to be added to the signal level at both ends. For simplicity in this analysis, the noise term is assumed to be identical in both devices. Solving this equation for the ratio of the distances in terms of the position-induced energy difference threshold yields the somewhat non-intuitive:

$$\frac{d_{NC-CRR}}{d_{T-T}} = \frac{\left(-\dfrac{d_{T-T}^{-2}\left(d_{T-T}^{2}N_0 - K_{path}P_{NC} + d_{T-T}^{2}N_0\left(\dfrac{S_{TCR}}{S_{RCR}}\right)\right)}{K_{path}P_{NC}\left(\dfrac{S_{TCR}}{S_{RCR}}\right)}\right)^{\frac{-1}{\alpha}}}{d_{T-T}} \qquad (5\text{-}3)$$

The different conditions are shown in Figure 5-5 for a signal (10 MHz bandwidth, 10 dBm at 2.0 GHz, -170 dBm/Hz noise floor, 0 dBm antennas, and an idealized "flat earth" propagation) at a range of 0.3 km between the non-cooperative transmitting node and the transmitting cognitive radio node. The *Range Ratio* in the *x* axis of the figure is the ratio of distance between the cognitive receiver and non-cognitive transmitter, compared to the distance between the cognitive and non-cooperative transmitters, determined by an equivalent expression to Equation (5-3):

$$\text{Range Ratio} = \left(\frac{d_{NC-CRR}}{d_{T-T}} \right) \tag{5-4}$$

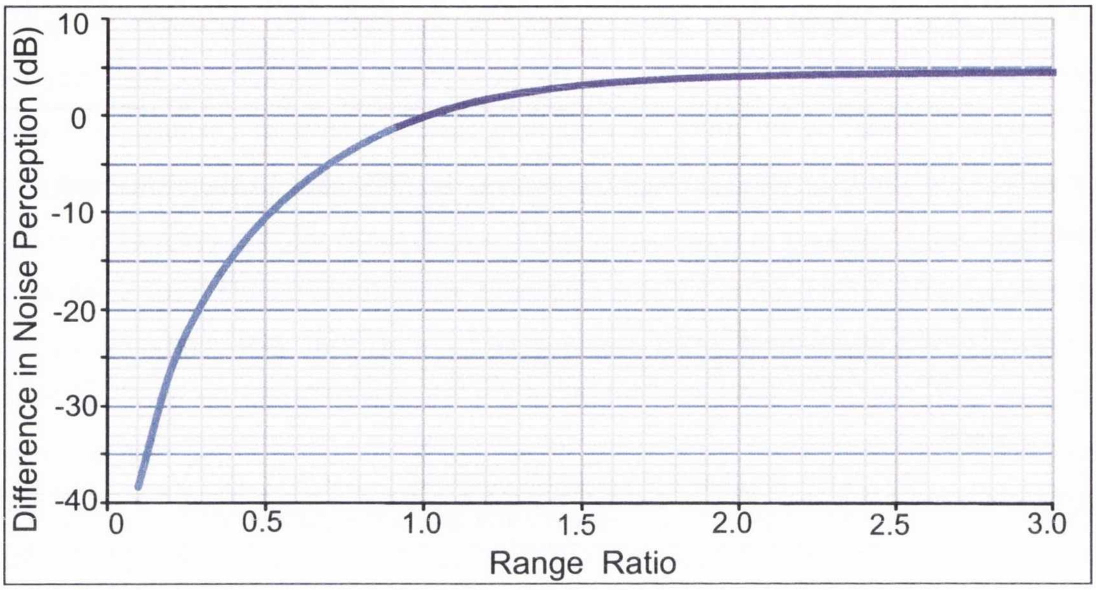

Figure 5-5 ***Illustrative Perception of non-Cooperative Signal on an ideal r^2 Plane***

Beyond a range ratio of 3, the signal is so weak that the entire difference in perception is the strength at the closer node (i.e., its level above the fixed noise). At equal distance the signal is perceived identically, and as the signal gets closer to the cognitive receiver (Range Ratio < 1), the perception of the signal rises rapidly at the closer node.

The same conditions can be considered, but with a policy prohibition that no frequency can be used if either end of the link perceives it at a threshold above the noise floor. In this case, the points with the strongest disagreement between the two nodes generally eliminates from consideration through the policy threshold rule. Figure 5-6 illustrates the same data shown in Figure 5-5 for a case in which the policy requires signals to be within 6 dB of the noise floor. The points precluded by policy are shown in Red, with an “X” mark over them, while the permitted points are in Blue. The range of signal amplitudes detected and permitted in policy controlled DSA operation (ranging from the noise floor to the occupancy threshold) is not significant compared to other local perturbations in the link environment, which can include fading, attenuation and other effects that are much more significant than the typical occupancy thresholds proposed.

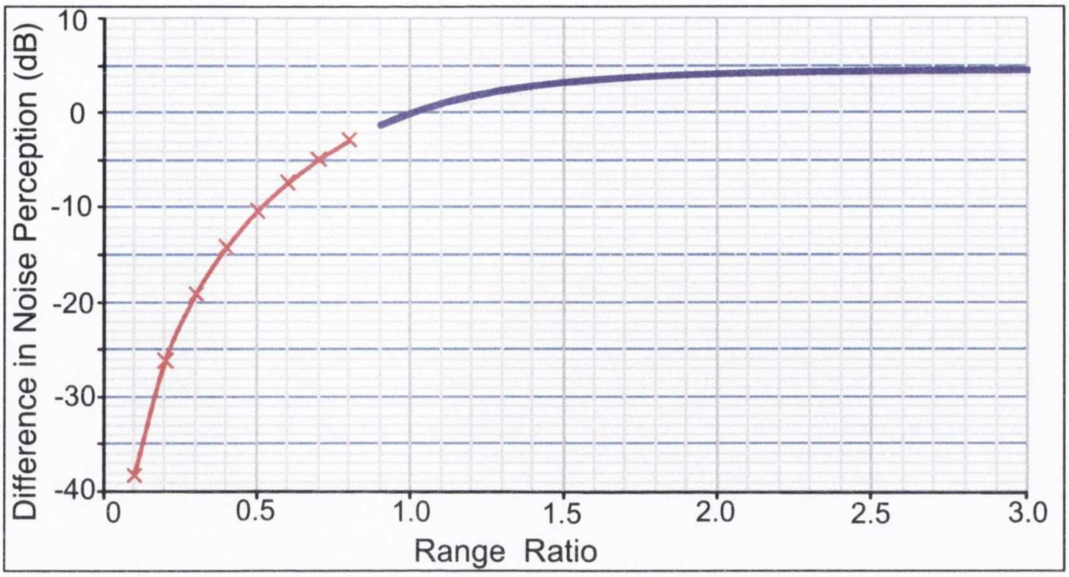

Figure 5-6 Illustrative Perception of Non-Cooperative Signal with Threshold Policy Controls

To determine the range of energy perceptions that would occur, the range of communications (using the same propagation exponent) is compared to the regions with significant difference (between cognitive radio link pairs) in detected interference energy level. In most environments, the regions with significant differences are quite small. In fact, in a link pair, the real risk is only half this level, since there is an even chance that the selection will be made by the node with the worst case interference, yielding the right decision regardless of the algorithm.

This situation is typical for the case of operation in an unlicensed band, when any available frequency can be used, regardless of the effect on any existing user of the frequency. In policy controlled spectrum an additional constraint is that neither node should detect an existing signal of more than *threshold* magnitude above a set noise level. The physical dimensions of the sensing regions are dependent on the specific quantities (frequency, energy, etc.) selected. However, a graphic depiction of these regions for a typical situation is instructive.

The next two figures illustrate a 10 by 10 km region containing two nodes that are spaced 2 km, and in line-of-sight with each other and interfering nodes. Figure 5-7 illustrates this case when there is no policy constraint on the signal level of available channels. In this case, the perception of the two nodes can be quite different. There is a possibility that a node could be close to one of the link partners and would receive significant interference, when at the same time, the link partner sensed only an insignificant interfering signal level. The difference in interfering signal strength in the white area is less than 1 dB between the two nodes, and in the blue areas, each color step is a 1 dB increase in signal strength difference, up to a 10 to 15 dB region. An extremely small region in the vicinity of each node exceeds this level (colored light blue through red). Clearly, a node with these operating rules cannot presume that if the frequency was usable at one location that it would be usable at the link partners!

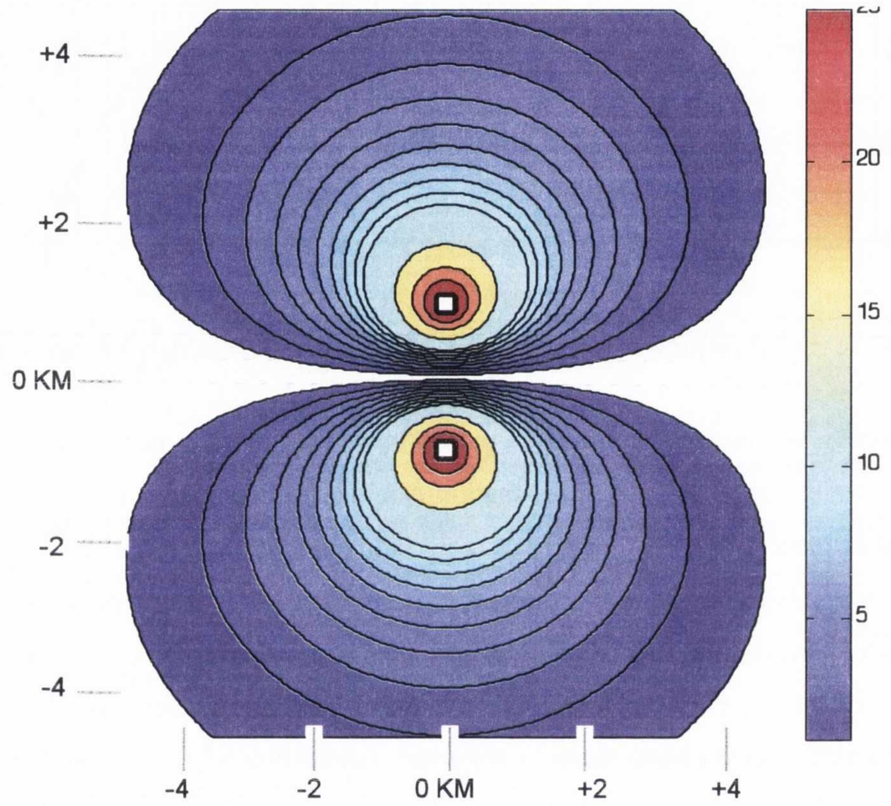

Figure 5-7 ***Illustrative Regions of Absolute Noise Floor Difference Between Two Nodes (in dB)***

Figure 5-8 illustrates the same situation, but with each node independently determining if the spectrum is occupied by applying a policy that the frequency is rejected if energy exceeds the threshold of 6 dB over natural noise floor.

If the interfering node is within the dark red region, both nodes detect that its power exceeds the threshold of 6 dB above the noise. In the orange region, one of the nodes detects that the threshold is exceeded. In the dark blue region, the signal levels are within one dB, and the lighter blue regions have increasing difference in energy, with a small “sliver” having a 3 to 4 dB difference. This region is not very significant compared with the other cases, and demonstrates there is only a limited benefit obtained with mutual determination of noise levels, when operating with an threshold policy. Any quantification is situation dependent, due to the number of independent variables in Equation (5-2).

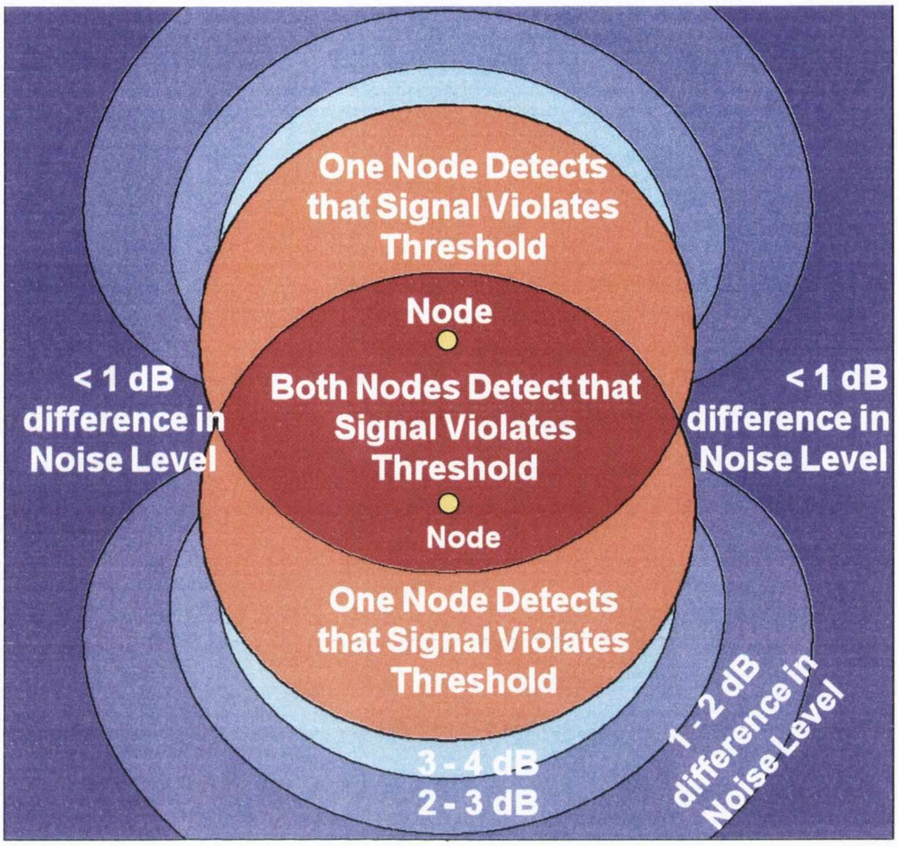

Figure 5-8 ***Illustrative Regions of Threshold-Constrained Noise Floor Difference Between Two Nodes***

Figure 5-9 and Table 5-3 provide examples that illustrate the statistical distribution of the various cases that arise in a radius from the midpoint of two nodes, out to a radius where the signal was 6 dB below the noise floor.

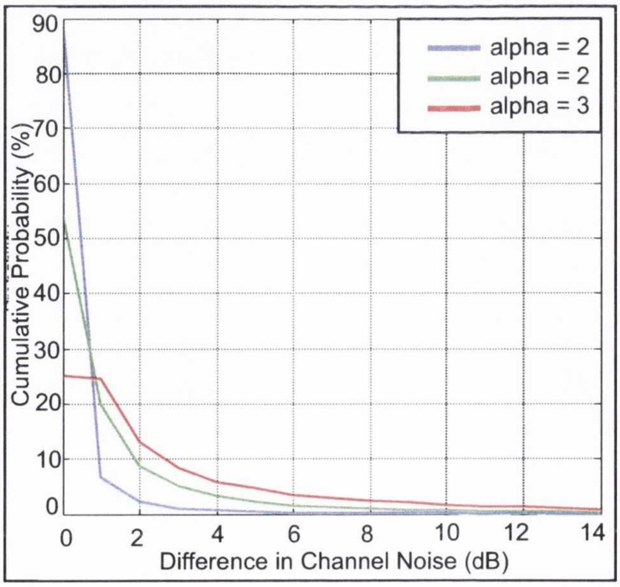

Figure 5-9 Statistical Distribution of Energy Sensing Differences without Policy Constraints

Table 5-3 Distribution of Interference Levels Among Link Partners without Policy

Node Difference	**% of Locations**		
	α =2	**α =3**	**α =4**
0 – 1 dB	88.7	54.1	24.9
1 – 2 dB	6.7	20.0	24.4
2 – 3 dB	2.1	8.9	12.9
3 – 4 dB	0.92	5.0	8.1
4 – 5 dB	0.54	3.3	5.7
5 – 6 dB	0.28	2.2	4.5
6 – 7 dB	0.18	1.4	3.3
7 – 8 dB	0.15	1.15	2.7
8 – 9 dB	0.1	.88	2.2
9 – 10 dB	0.05	.64	1.9
10 – 20 dB	<.02	<2.1	6.4
20 – 30 dB	<.01	<.5	1.6
30 + dB	0	<.5	<1

These distances were scaled to achieve this energy at the limits of the operating area for a range of propagation characteristics (α from 2 to 4). Node spacing was not varied, and was held constant at 1 km. Not surprisingly, as the propagation exponent increases, the spectrum environment becomes more local, and it is increasingly likely that there will be significant differences in spectrum perception between the nodes.

Figure 5-10 and Table 5-4, illustrate the case where a spectrum access policy required both nodes to determine there was no signal in the channel more than 6 dB above the noise floor. This has the effect of ensuring that the difference between nodes cannot exceed the threshold value. [Note that the percentages in the graphic are the percentage of policy permitted operations, and do not include the policy-precluded cases shown in Table 5-4.]

Table 5-4 Distribution of Interference Levels Among Link Partners with Policy

Difference at Node	**% of Locations**		
	α =2	**α = 3**	**α = 4**
Fails Policy @ one Node	3.2	10.1	20.1
Fails Policy @ Both Nodes	6.8	13.5	18.9
0 – 1 dB	86	52.3	23.1
1 – 2 dB	3.2	17.7	22.6
2 – 3 dB	0	5.5	10.3
3 – 4 dB	0	0.13	4.75
4 – 5 dB	0	0	0.15
5 – 6 dB	0	0	0
6 + dB	Not Possible		

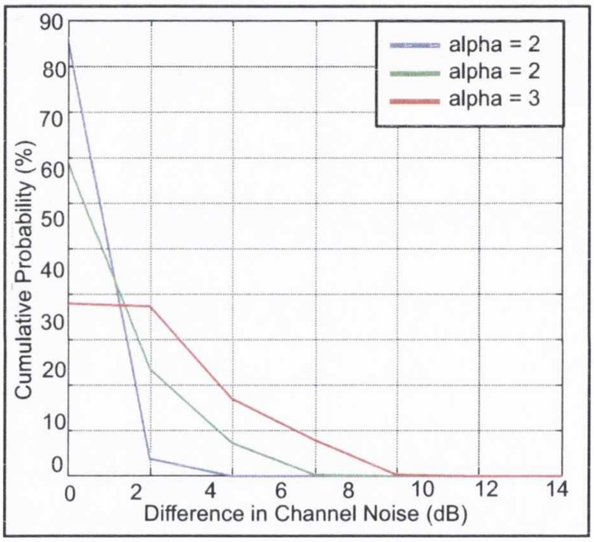

Figure 5-10 Statistical Distribution of Energy Sensing Differences with Policy Constraints

Inspection of both the closed-form expression and the empirical results demonstrate three fundamental principles:

1. The enforcement of a policy threshold jointly by link partners has the effect of inherently minimizing the noise floor differences between the nodes of the link.
2. The propagation characteristics are a significant driver in how much difference can appear between link partners. This is true because the range over which reliable communications can be performed is significantly less than the range over which interference can be caused through influencing the noise floor. Nodes operating in short, line of sight paths inherently will experience high degrees of channel noise floor correlation, while those operating over diffracted paths will have a significant possibility of differences in noise floor at different ends of the path.
3. In policy controlled spectrum, with low thresholds above the noise floor, the advantages of multilateral determination of operating frequency is generally not significant.

Table 5-5 illustrates the performance implications of frequency selection (in terms of noise floor) by either endpoint, or when the decision is made jointly.

Table 5-5 Effectiveness of Noise Floor Minimization Algorithms[1]

		Transmitter Selects	**Receiver Selects**	**Jointly Select**
Algorithm		The transmitter selects a (policy compliant) frequency that has the lowest signaling channel noise present. Receiver refuses to rendezvous if it locally violates threshold policy.	The receiver selects the channel as part of the rendezvous process. It may fail to rendezvous on transmitter selected channels, or direct a change to a preferred frequency.	Both nodes search through all candidate frequencies to determine the one with the least mutual impact.
Link Suitability		Simplex / Duplex	Simplex	Duplex
Within Fixed Thresholds	**Weaknesses**	Transmitter selection could be several dB from the receiver's environment.	More complex and time consuming rendezvous.	Most complex and time consuming rendezvous. May thrash in dynamic environment.
	Benefits	Simplest algorithm, as all decisions are made at one node at one time.	Can minimize noise floor. Potential benefit is only several dB.	Can address highly varied diffracted environments. Potential benefit is only several dB.
No Threshold Constraints	**Weakness**	Slight chance that there is a significantly different environment (>10 dB) at receiver.	More complex and time consuming rendezvous.	Most complex and time consuming rendezvous. May thrash in dynamic environment.
	Benefits	Simplest Algorithm, as all decisions are made at one node at one time.	Slight chance of a significantly lower noise floor at receiver (>10 dB).	Can address significantly different environments at each end.

The transmitter selects first in the base case, since this is the simplest to implement in a DSA scheme.

When operating without fixed thresholds, the potential conditions are much more variable, as it would be policy permissible to operate on a channel when one node was in close proximity to a signal that might not even be detectable by the link partner. This situation is typical of Wi-Fi installations in the unlicensed bands where there is no requirement to

---

[1] This discussion is independent of the specific mechanism used to perform the rendezvous operation, which is a significant complexity in the implementation of a cognitive radio. The discussion addresses the decision to rendezvous the nodes on a particular frequency, not the process used to accomplish it.

avoid interference to other users, and therefore communications may be initiated with very asymmetric noise distributions.

In this table, the "transmitter selects" and "receiver selects" refer to the option to make the first choice of frequencies. The other partner can veto a choice by refusing to rendezvous, which would be the case if the channel violated its policy, but cannot offer an alternative that is more optimal if the initially selected frequency is non-optimal, but otherwise permitted by policy. In "Jointly Selects", the two nodes perform a decision process that reflects both the policy compliance at both ends, and a joint evaluation of optimality.

5.4 Noise Floor Management Benefits

The benefits of noise floor selection are essentially the difference between the best and the likely noise floor within permissible frequencies of operation. For DSA in highly regulated bands, the interference policies may well preclude operation on any occupied frequency, capping potential gain at the difference between the thermal noise floor and the maximum permitted in-band power. For an unlicensed band, where operation on any frequency was permitted, the benefit would be quite high, depending on the average spectral noise levels above the thermal nose level. Rudimentary versions of this algorithm are in fact often included in devices intended for operation in shared bands as a form of automatic channel selection.

5.5 Noise Floor Management Conclusions

Complex noise floor management algorithms are of only marginal value when operation of nodes is constrained by tight interference policy thresholds. The individual and joint perceptions of the channel energy by two nodes must be very close to the noise floor in order for the frequency to be mutually acceptable under a threshold policy. The distribution is very close to the noise floor for threshold compliant signals. The probability of a significant difference in noise floor, and that the frequency is compliant at both ends of the link is reasonably low.

When there is no threshold enforced, the benefits of joint selection of a channel are much more significant, particularly when the potential interferer could be close to one link partner and undetectable by the farthest node. In the absence of a joint implementation of a threshold policy, it is essential that both nodes collaborate in the selection of channels.

Chapter 6 Minimization of Interference Effects through Interference Tolerant DSA Mechanisms

6.1 Dynamic Spectrum Access Role in Interference Avoidance and Tolerance

The DSA mechanism is an enabling requirement for many of the techniques and technologies discussed in other chapters (4 thru 8), and will likely be an inherent element in most adaptive wireless structures. The DSA-enabled techniques include front-end linearity management, dynamic topology, and waveform selection. Even if the spectrum management benefits were non-existent, the flexibility DSA affords to manage other aspects of the wireless environment justifies its inclusion in wireless devices. DSA offers capability benefits through pooling[1] of spectrum resources, and managing spectrum conflict resolution dynamically. It is a building block in the transition from existing spectrum practices to an "ecosystem" of cognitive radios. Additionally, analysis of the bounding capability of wireless provided by Gupta and Kumar in their seminal paper demonstrate that interference is the constraint on wireless network scalability [141]. Implicitly, mitigation of interference is critical to achieving, or exceeding this bound.

DSA research has been premised on the belief that significant benefits can be achieved through sharing of spectrum with incumbent systems. In this DSA concept, spectrum is shared, but absolute protection is afforded to a specific protected class of users. In this chapter, it will be argued that even more significant gains can be achieved when wireless operation transitions to spectrum sharing between dynamic and interference tolerant systems, where absolute protection is provided to no node, and service is assured statistically. The deployment of cognitive radio will deliver two increments of benefit; the first, when cognitive radios more effectively share spectrum with non-cognitive radios, and second, when cognitive radios can assume that other radios are cognitive, and can assume that they are capable of mitigating any interference situations that do occur.

There is a significant literature that has described opportunities, approaches and implementations for spectrum sharing under principles of non-interference to the incumbents ([28], [65], [78], [140], and [142] for example). This chapter will also address operating regimes in which all occupants of a given spectrum allocation can assume that other devices are capable of mitigation and tolerance of some degree of interference.

---

[1] Spectrum pooling is defined as aggregating spectrum for multiple systems, applications, users or other community participants into a single set of permitted frequencies, for which all members can contend to access spectrum.

While it is not desirable to cause interference, this permits all occupants of the spectrum to operate in dense deployments that cannot be supported with interference-free spectrum assignments. There will be higher risks of causing and receiving interference compared to situations in which high confidence of non-interference must be guaranteed. This chapter will demonstrate that interference tolerant operation has significant benefits to all of the devices in the spectrum. They can provide significantly more density and aggregate capacity than either manual spectrum deconfliction or primary/secondary non-interfering DSA operation, and while still offering arbitrarily high levels of reliability through statistical and adaptive mechanisms.

Frequency management regimes are classified as follows:

Manual Frequency Planning	Manual deconfliction of frequency assignments to ensure non-interfering operation for all users and potential interferer operating areas. This is the situation when non-cognitive radios are assigned frequencies in the vicinity of other non-cognitive radios.
Interference Free Secondary Sharing	Automated, or preplanned sharing of spectrum under principles that ensure that secondary users have essentially zero probability of interfering with primary usage. Secondary users cannot assume access to specific frequencies or spectrum. This is the situation when cognitive radios are sharing spectrum with protected, non-cognitive radios.
Interference Tolerant Sharing	Interference avoidance through bilateral sensing of the environment and mutual reduction of interference, with all nodes responsible for mitigating the effects of the interference they receive, and balancing the interference they might cause. This could be the situation when cognitive radios share spectrum with other cognitive radios.

Much of the current DSA research is focused on the second alternative: Interference Free Secondary Sharing. This approach to spectrum sharing is designed primarily to rely on physical layer techniques to minimize energy in the communications channel. However, to create interference tolerance, it is necessary to share the burden of operating in interfering environments through multiple levels of levels of the communications architecture. In interference tolerant operation, the node must be able to mitigate the effects of temporary interruptions to service, and also ensure that its own operations do not preempt other users' access to the spectrum to such an extent that it causes highly correlated (long duration) outages to other nodes. In effect, these operational mitigations attempt to decouple interference events that occur to the physical channel from the reliability of the overall service being provided. Centralized control over these decisions has been proposed by Kovacs and Vidacs [143]. A quantitative analysis of the probability of interference of such regimes is provided by Win, in which it is assumed that a secondary user's failure to detect a primary user results in interference [144]. This is appropriate for the special case in which cognitive radios share with non-cognitive ones, but an interesting, and more robust regime is established when all users of the spectrum are capable of both dynamically locating spectrum, and of determining that their operation is being impacted by interference, and relocate themselves automatically.

The intent of this section is to analyze approaches that are capable of being autonomously implemented within a peer-based structure. Some of these operating features required to transition from interference avoidance to interference tolerant are shown in Figure 6-1.

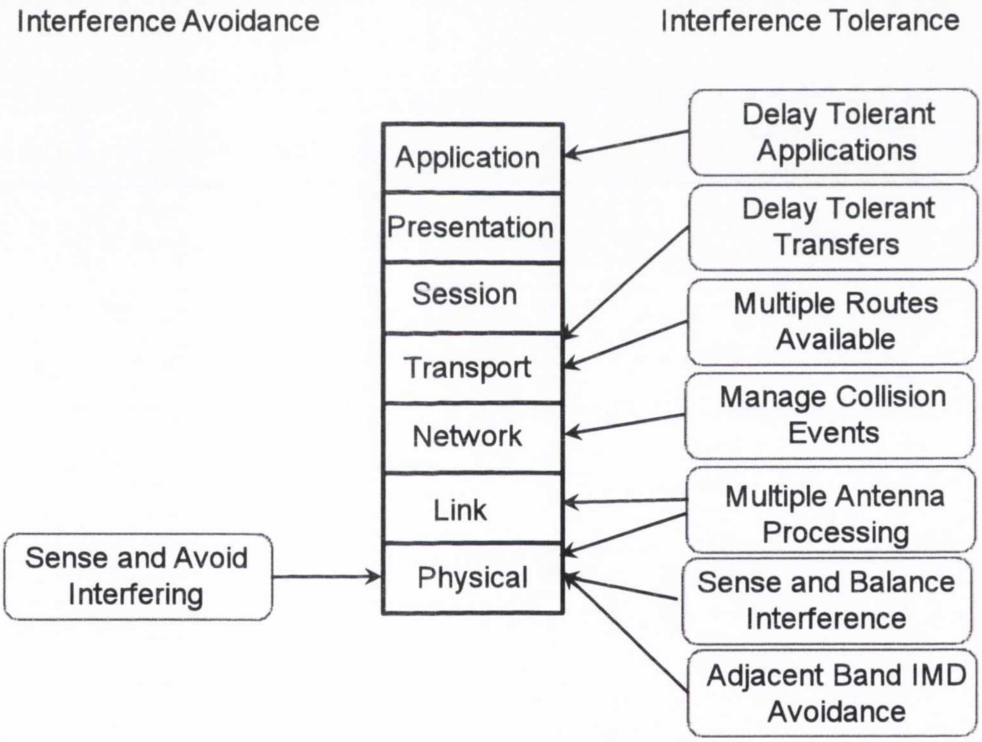

Figure 6-1 Interference Mitigation throughout the ISO Communications Layers

This approach is the opposite of current practice, in which the PHY layer had the total responsibility for interference free channel operation, with an implicit guarantee of service to upper layers. In the approach illustrated above, addressing the consequence of interference is an integrated and shared responsibility of all layers. By mitigating the effect of interference, the resulting aggregate network capability (summation of all node throughput) can be greater in terms of both capacity and reliability. Specific layer responsibilities are shown in Table 6-1.

Table 6-1 Interference Mitigation Techniques in ISO Layers

Feature	Layer	Description
Adjacent Band IMD Avoidance	PHY	Dynamically select operating frequencies with an awareness of filter performance, total energy, and IIP3 characteristics.
Interference Balancing Physical Layer	PHY	Determine the effect of channel selection on other users and self to balance interference across networks and users. In effect, this is Gaussian "whitening" of the interference.
Multiple Antenna Processing	PHY/ MAC	Processing multiple antenna returns of a single signal to isolate interference and improves SNR, such as in 1xN Multiple Input/Multiple Output (MIMO) operations.
Collision Detecting MAC Layer	MAC	Detect when the channel is busy by either other network members or other users of the frequency before attempting to transmit.
Multiple operating Links	NET	Provide the network multiple paths to and nodes so interruption of any one path does not impact high-QOS traffic. As examples, this could be implemented as one transceiver time sliced between frequencies, or multiple transceivers within the node.
Delay Tolerant Networking	Transport	Organize end-to-end transfers on a hop-by-hop basis so delays or losses in one hop do not cause information transfers to be lost, reset, or initiating/receiving application to fail [85].
Delay Tolerance	App	Applications do not make assumptions about network latency, except where critical to the application[1] [145].

As examples of this approach, an interference tolerant MAC layer could preferably use some variant of collision sensing multiple access (CSMA), in order that it could deal with uncoordinated activity on the same channel, without reliance on either cooperative mechanisms, such as RTS/CTS or deconfliction, such as agreed to slot allocations in Time Division Multiple Access (TDMA). In terms of waveform design, a cognitive radio would adapt its waveform strategies between spreading techniques, which would raise noise floor continuously, and short, higher energy usage intervals that would deny effective use of the channel for a period, but would not impact it for the free periods. Such decisions are inherently situational.

The choice of interference regime is also impacted by the nature of the spectrum usage. Interference free usage is an appropriate model for applications that make essentially continuous usage of the spectrum, such as analog, high duty cycle or broadcast, or when the receiver can not participate in adapting the link. Where transmissions are intermittent,

---

[1] Integration of applications within the framework of a cognitive radio is contemplated by cross layer approaches, such as outlined by Marojevic et al.

and the receiver can be involved, the interference tolerant mode of operation is particularly attractive. An analysis of these two treatments is provided by Sadek et al. [146].

Consideration of the relative interference levels considered acceptable in Physical and MAC layers is instructive. In creating interference free operation, most spectrum management processes attempt to create a very low probability interference incidents, which to be unnoticeable are typically constrained to occur at rates that are below 10^{-5} to 10^{-6} (interference events < minutes per year). By contrast, MAC layers operate at collision rates often above 10%, in order to ensure maximum utilization of the channel [147]. For example, Jardosh et al. report the collision statistics as the throughput of a WLAN is increased [148]. This different treatment of interference/collisions is highly unbalanced, and perhaps reflects historical approaches, rather than an engineering optimization across these processes. In this chapter, the benefits of allowing co-channel interference similar to MAC layer collisions will be examined and quantified.

The analysis of DSA will develop generalized equations for the physical separation of nodes (or density/area) as a baseline for the improvement in aggregate network throughput and spectral usage through DSA. The effect of fade margin, propagation exponent and duty cycle will be characterized in terms of non-interfering operation of DSA systems. The performance of DSA systems that accept and mitigate interference will be characterized to quantify the benefits of heterogeneous DSA systems, sharing spectrum to optimize performance in environments where radios both can cause, and must accept and mitigate interference.

Previous work has demonstrated that ad-hoc networks operating on single frequencies have severe interference limitations [141]. In examining spectrum utilization for cognitive radios, four spectrum usage and density situations will be considered, and generalized occupancy formulas for spectrum density will be developed.

$SD_{\text{NC-NC}}$	The density of non-cognitive, stationary, nodes that can be guaranteed a high level of interference protection, since they cannot be adaptive and mitigate in-channel interference for nodes whose location is known.
$\text{SD}_{\text{MNC-MNC}}$	As above, but for nodes that are also mobile within an operating region, which need protection throughout the region, and could cause interference to other nodes throughout the region.
$SD_{\text{CR-NC}}$	The density of cognitive radios sharing spectrum with non-cognitive radios that must be guaranteed a high level of interference protection, since the non-cognitive radios cannot be adaptive and mitigate external caused, in-

channel interference. No spectrum protection guarantees are assumed, or required, for the cognitive radios sharing these bands.

$SD_{CR\text{-}CR}$	The density of cognitive radios sharing spectrum with other cognitive radios that are not guaranteed interference protection, since the cognitive radios can adaptively mitigate in channel interference. The ceiling on capacity is given by maximizing the aggregate effective throughput of the radios in the band, rather than the performance of any given radio.

Before discussing a cognitive spectrum process, the classical spectrum management and assignment case is considered. Once radios moved beyond spark gap techniques (the original impulsive Ultra Wide Band (UWB) radio!), use of the spectrum has been de-conflicted to avoid interference. Spectrum and frequency managers assign individual radios or networks discrete frequencies, and attempt to ensure that the emissions from one do not adversely impact others. A not insignificant legal (and seemingly smaller technical) community has grown around this simple principle.

This chapter considers spectrum strategies that use awareness to locate spectrum holes that are themselves often the result of the essential conservative nature of the spectrum planning process. However, an equally important rationale for their inclusion in wireless systems is the ability to locally resolve interference by using the same behaviors used to locate new and unblocked spectrum. This feature of interference adaptive radios offers all users of the spectrum the ability to "back off" the current conservative assumptions that underlay spectrum planning.

Many approaches to spectrum planning have attempted to reduce the risk of interference. The approach outlined in this chapter instead focuses on an alternative strategy: reduce the consequence of interference, and thereby enable all devices in the spectrum to adopt less conservative assumptions, with the confidence that interfering conditions can be mitigated with tolerable consequences.

Another view of this process is to consider how the different frequency deconfliction processes resolve uncertainly in the characteristics of the devices sharing spectrum, planned versus actual propagation conditions, and the usage of the spectrum. A summary of how each technique addresses uncertainty is shown in Table 6-2. Each successive approach is able to more effectively resolve uncertainty in the operation of the links, although at the cost of introducing additional risks of interference. However, even when using DSA to share spectrum with non-interference tolerant devices, the algorithms must assume many of the worst-case assumptions on propagation due to the unacceptable risk of interference.

Table 6-2 Uncertainty Management in Frequency Management

Uncertain Characteristic	**Manual Frequency Planning**	**Non-Interfering DSA**	**Interference Tolerant DSA**
Temporal Usage Patters	Typically not exploitable, and channels are reserved as if used 100% of the time	Senses Channel Usage, and only considers interference if no activity is detected	Same as non-interfering DSA, but may choose to share and accept interference
Mobility Patterns	Spectrum deconflicted for all possible locations of primary users	Implicitly senses interference potential of actual location	Same as non-interfering DSA
Propagation between Victim Transceivers	Assumes poorest propagation, which is high propagation loss	Assumes poorest propagation, which is high propagation loss	Determines risk based on likely distribution of propagation loss
Propagation from Victim Transmitter (to DSA node)	N/A	Assumes poorest propagation from source to sensor	Determines risk based on likely distribution of propagation loss
Propagation to Victim Receiver (from DSA node)	N/A	Assumes signal will maximally propagate from source to sensor	Determines risk based on likely distribution of propagation loss

Another motivation to consider transition to interference tolerant operations is the bounds provided by Gupta and Kumar in their analysis of the throughput of wireless networks [141]. In this work, they demonstrate that as the network density increases, the effect is to have the interference level created become the dominant factor in aggregate network throughput. The performance of wireless systems then must recognize that in the end, spectrum is finite, and scalable performance will depend not on operating capability in clear spectrum, but in the ability to maximally exploit dense and interfering environments through adaptations in all layers of the wireless system.

An additional consideration is that the responsibility to tolerate interference can be supplemented with responsibilities that maximize the probability of being detected. If a device uses techniques that are difficult to detect, than that device must expect and tolerate additional interference that results from the lack of detection by other users of the spectrum. Friend and MacKenzie demonstrate that specific topology choices can greatly decrease the probability of a secondary network causing harmful interference [149].

In the following discussion, it is assumed that a cognitive radio has the capability to:

1 Sense the local spectrum

2 Compare spectrum measurements with profiles of permitted usage

3 Adjust its frequency of operation

4 Make reasoned assumptions about, or measurements of, the range of propagation/attenuation conditions

Proposed spectrum sharing approaches have included ultra-wideband (UWB) approaches, which share all frequencies at a low energy level, and narrowband, which have a single user at any given time and place. A hybrid of these is the "interference noise temperature" approach introduced in the FCC Spectrum Policy Task Force Report [136]. While the noise temperature approach has attracted significant research interest [150]-[154] it does not appear likely that it will be adopted as a national policy in the near or medium future. Therefore, the term DSA throughout this document will generally imply dedicated usage of segments of spectrum, in at least a local region. The general principles developed are extensible to the wideband spreading of energy as contemplated by the interference temperature concept.

6.2 Analysis Approach and Assumptions

There are four important considerations in examining the effect of various approaches to spectrum management. These are:

1. Communications Range and Receiver Characteristics
2. Mobility
3. Propagation Characteristics
4. Usage and Operating Characteristics

These are further described and quantified in the following sections.

6.2.1 Communications Range and Receiver Characteristics

Spectrum and frequency planners are inherently disadvantaged by a number of factors. For one, they have to assume that:

- Interfering signals will propagate to the maximum possible range (which reduces power or increases standoff distances).
- Desired signals will need to be received without unacceptable link margin degradation in the worst-case propagation conditions (which increases required power, and lowers noise tolerance).

In practice, this means that interference analysis is often driven by two unlikely conditions, maximal propagation of interfering signals, and minimal propagation of the desired signal. Although individual situations vary widely, the likely range of conditions has been measured, and a number of environments have had their distribution characterized [155]. To summarize the measurements, the propagation exponent (n) had a mean of 2.7, and a variance of 0.6.

Figure 6-2 illustrates the relationship between the various communications and interference ranges involved in DSA. Case (a) is the desired high assurance communications range, which must assume worst case propagation (α_{wc}) and fade, and still assure a signal level above $E_{receive}$. This is a conservative range, but typical for the assumptions required for high assurance link planning. Case (b) is the distance by which radios must be separated for manual frequency deconfliction, reflecting that the victim radio may be in an advantaged position to the DSA transmitter (α_{bc}), with no fade condition present, and that the signal level in that situation must be below the interference threshold ($E_{interfere}$). The DSA separation is shown in Case (c), between the best and worst case propagation (α_{tc}), significant fading is not present, and the energy interference is also limited to $E_{interfere}$. In the discussion in this chapter, it is assumed that the bandwidths in use are comparable.

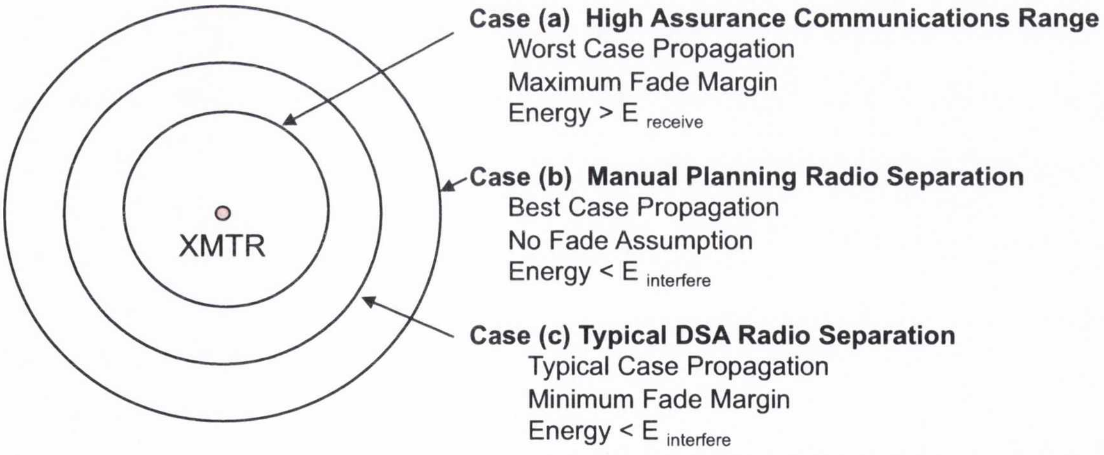

Figure 6-2 Practical Interference Margins

Most of this pessimism is not inherent in the operation of the radios themselves; it is strictly a consequence of having to plan for "edge cases" in advance of knowledge of the actual conditions. Static spectrum planning must assume that links operate at the maximally stressing conditions, while interference will occur when links are maximally aligned and propagating to cause interference. DSA enables cognitive radios to reduce both of these risk margins to reflect the actual conditions present for both intended and victim links. This enables increased density through two mechanisms:

(1) Cognitive radios can exploit spectrum by adapting to other spectrum users' actual usage (time, frequency, power, etc.) and thus, create more assignments in the same spectrum.

(2) Cognitive radios can tolerate more aggressive spectrum reuse because they can move their spectrum assignments in response to any interference they do receive, allowing other radios to be less conservative in sharing spectrum with them. Table 6-3 provides the model variables that reflect range and receiver characteristics.

Table 6-3 Communications Range and Receiver Characteristics

R_{wc}	Worst Case Range	Unit	Required range corresponding to the maximum fade and the worst case α
$P_{receive}$	Required Receive Energy	dBm	Required energy at the receiver in the worst case range condition
$E_{interference}$	Interference Threshold	dBm	Maximum energy present on a channel in order for that channel to be reused

6.2.2 Mobility Characteristics

As stated previously, the mobility metric is driven by the requirement that spectrum be de-conflicted out to the interference range from the entire mobility range where the device might be located, as depicted in Figure 6-3.

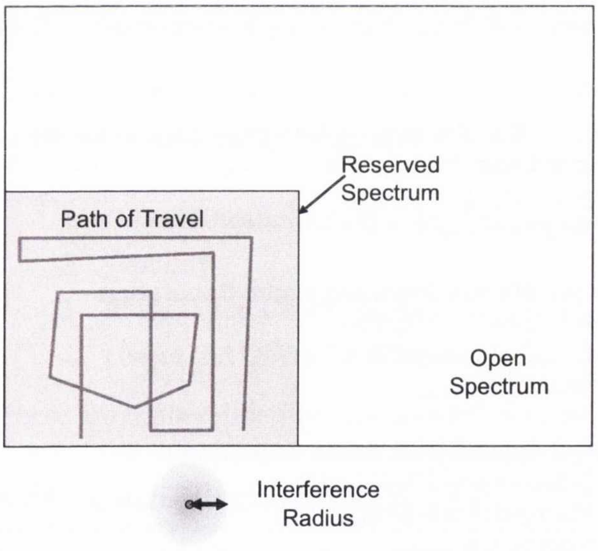

Figure 6-3 Determination of Mobility Factor in Spectrum Management

In this figure, the spectrum manager has awareness of all of the paths (or locations of operation) where the device might be present, plus the interference radius by which it must be de-conflicted. This creates a reserved area (actually a volume, but the third dimension is rarely exploitable, so is ignored for this discussion) within which spectrum cannot be used for any other purpose without risk of mutual interference. If there is no pre-knowledge of mobility, the spectrum must be reserved over the entire extent for which operation is permitted. Note that interference radius must be bilateral. It must include the effect of the power spectral density of both devices. In determining the interference radius, it must reflect the larger of the two directionalities; i.e., the greater of the range at which *A* might interfere with *B*, or B might interfere with *A*.

Table 6-4 illustrates the model variable that reflects mobility characteristics.

Table 6-4 Dynamic Spectrum Mobility Characteristics

Area	Operating Region	Unit^2	2 dimensional area over which nodes operate, in the same units as R_{wc}.

6.2.3 Propagation Characteristics

Actual link characteristics are unknown to the link planner, human spectrum manager, or DSA component of the radio. In the following discussion, power management enables the link to operate with power appropriate to the actual conditions of the link, rather than those of the worst case conditions and fade margin. In many networks it is impractical to provide power management. The hardware used in the network is often not capable of significant ranges of output power; the feedback mechanisms cannot be provided in a simplex device; there is a significant number of receivers that must receive the broadcast; or some of the receivers have no feedback mechanism at all, such as broadcast networks with receive only terminals. Although power management may be present, the spectrum planning process must assume that it does not reduce the maximal radiated power.

Link and interference propagation characteristics that are reflected in density modeling are provided in Table 6-5.

Table 6-5 Dynamic Spectrum Propagation Characteristics

F_{margin}	Fade Margin	Ratio	Degree of additional fade margin added to the link budget in the worst-case range condition. 1 = no fading, 10^{-2} = 20 dB margin.
A_{wc}	Worst-Case Propagation		Propagation coefficient corresponding to the worst-case conditions ($r^{\alpha wc}$).
A_{bc}	Best-Case Propagation[40]		Propagation coefficient corresponding to the best-case conditions ($r^{\alpha bc}$).
A_{tc}	Typical Propagation		Propagation coefficient corresponding to the typical conditions ($r^{\alpha tc}$).

6.2.4 Operating Characteristics

The *duty cycle* is the amount of time for which spectrum must be reserved compared to the actual time during which it is actually used. The definition of the term "*used*" is not necessarily just the time for which energy is being transmitted. Its evaluation is somewhat subtle due to two effects. First, any time in which the receiver is sensitive to interference is certainly use of the spectrum, and should be considered as "used". Transmitting on a

---

[40] The units of distance for propagation are arbitrary, but should be selected such that they are similar to the distance at which the best and worst-case propagation diverge; i.e., the point at which propagation transitions from direct to diffraction. Typically, this might be 200 to 500 meters for a low antenna, 1 km for a handheld, and would be irrelevant for a space link.

beacon channel could be disruptive, even if there is no active transmitter. Secondly, time periods that are too short to be exploited by other users are essentially "used." Interleaving independent (non-cooperative) users within MAC layer intervals is generally not practical with current MAC layers, so the entire operating period should be considered as being occupied, since the brief openings are not practically exploitable. The operating characteristics and spectrum assumptions are provided in Table 6-6.

Table 6-6 Dynamic Spectrum Operating Characteristics

Duty	Duty Cycle	Ratio	Ratio of time transmitting on a given network to the total time.
$N_{\text{Recipients}}$	Number of Units	Count	Number of Units that must receive a broadcast.
N_{Pool}	Pool Size	Count	Number of channels that are contended for among the spectrum users.
A_0	Availability	Ratio	The availability that must be provided for immediate access to the spectrum.

6.2.5 Analysis Approach

A postulate is that cognitive radio offers the ability to manage a spectrum situation more effectively by utilizing the ability to sense the actual propagation conditions that occur, and to adjust the radio dynamically to best fit these conditions. Its operation is considered in the pursuit of two objectives. First, it attempts to minimize its own spectral "footprint," consistent with the environment and needs of the networks it supports. Second, it adapts itself to fit within whatever spectrum is available, based on local spectrum sensing. When these two principles are integrated, they conceive a radio that can both locate spectral openings, and morph its emissions to fit within one or more of these opportunities. Such radios could offer radio services without any explicit assignment of spectrum, and still be capable of providing high confidence services. It is assumed that the interference policy precludes selection of frequencies that would be disruptive to other protected spectrum uses.

For purposes of analysis, the effectiveness of a long existing communications protocol; generally referred to as "*Listen Before Talk*" (LBT) will be considered. Radio operators have used this technique to avoid interference since the first days of radio communications. Leu [156], McHenry et al. [78], [157] and Seelig [158] describe its adaptation to operation in DSA radios. There are potential weaknesses in LBT, but most of them lead to, or cause interference to other wireless networks, rather than the cooperative cognitive radio, therefore they are not significant from the perspective of cognitive radio performance. However, these weaknesses are of critical importance to the regulatory community. Other research, such as that sponsored by DARPA's XG project has published analysis of these algorithms, and their performance, such as McHenry [140], [159] and Marshall [142].

6.3 General Approach for Non-Interfering Operation

One fundamental difference in performance between cognitive and non-cognitive radios is in how they obtain and access spectrum. Non-cognitive radios generally obtain spectrum in one of two ways:

Assigned:	Assigned spectrum is typically guaranteed to a given user or usage from a regulatory authority (or by delegation from one) and is typically assumed to be exclusive or preemptive use. Typically, there is an assurance of non-interference with this class of spectrum. Broadcast services, cellular, satellite, and public safety are examples of this class.
Commons:	Commons spectrum is provided for use by a number of users. Generally there are some technical or operational constraints. There is no assurance of availability or non-interference. Examples of this class include the Industrial, Scientific and Manufacturing (ISM) bands, commonly referred to as unlicensed.

This dissertation examines the problem of spectrum assignment using manual planning and deconfliction as the reference point for comparison. This mechanism is the current practice, and is an appropriate baseline. From the perspective of a communications link, there is nothing that is preferential to a perfectly managed spectrum band and spectrum that is sufficient for, and dedicated to, its operation. In practice, dedicated spectrum must be a compromise between the amount of spectrum dedicated solely to individual uses, and the requirement to maintain interference free operation. A wider channel allows for more capacity, but reduces the total channels available to be allocated to users and services, and therefore increases the potential for interference. Chapter 7 will examine another aspect of spectrum assignment and waveform selection, which is the benefit of dynamic signal footprint management.

The performance impact from introduction of DSA is an aggregate function that reflects not only the impact on the radio's operation, but also the collective effect of the radio on all other devices sharing the spectrum (as pointed out in Chapter 4, this is not limited to the signaling frequency, although this chapter will focus on the effects within the signaling bandwidth). To determine these effects, the practical constraints and limitations of non-DSA spectrum management of the devices in the spectrum need to be considered.

One of the underlying principles of DSA is that dynamic deconfliction can space radios at closer distances than manual planning in both spatial and temporal dimensions. This ratio can be analyzed by determining the lowest possible radiated power[41] (P_T) and path loss constant (K_{path}) required to reliability operate a link, in terms of the link parameters shown in Table 6-3 thru Table 6-6.

Since, by definition, for the worst case link:

$$\frac{P_T K_{\text{path}} F_{\text{margin}}}{R_{wc}^{\alpha wc}} = E_{\text{receive}} \tag{6-1}$$

where:

P_T is the transmitter power
K_{path} is the loss over a unit distance
all other variables are as defined in Table 6-3- Table 6-6

The manual frequency manager must assume that interference will propagate at the "best case" propagation constant, and without fade, so the range to the closest receiver reusing the frequency (S_{manual}) must be at least:

$$\frac{P_T K_{\text{path}}}{S_{manual}^{\alpha bc}} = E_{\text{interfere}} \tag{6-2}$$

where:

S_{manual} is the separation distance required for worst case, manual spectrum deconfliction

This is necessary because manual planning must consider that a source of interference and a recipient of it may both be in an advantaged position relative to each other, such as on opposite hills or mountain tops and have r^2 propagation (α_{bc}), while the intended recipient is over a hill and in r^4 (α_{wc}) conditions.

$E_{\text{interfere}}$ reflects the maximal signal level permitted in comparison to the weakest possible exploitable receive signal level. From this, E_{recieve} can be related to $E_{\text{interfere}}$ (Signal to Interference and Noise Ratio) and the receive range can be rewritten in terms of the interference level permitted. Assuming required signal is much greater than noise:

$$SINR = \frac{Signal}{Interference} = \frac{E_{\text{receive}}}{E_{\text{interference}}}, \qquad E_{\text{receive}} = \text{SINR}\ E_{\text{interfere}} \tag{6-3}$$

$$\frac{P_T K_{\text{path}} F_{\text{margin}}}{\text{SINR}\ R_{wc}^{\alpha wc}} = E_{\text{interfere}} \tag{6-4}$$

---

[41] P_T here is used as a non-directional EIRP. The same formulation could express P_T as a function of an angle θ, in which case the spatial relationships of the separation distances would also become functions of θ, and would be unique for each θ value.

The assumption that maximum interference energy is related to minimum receive energy lets the separation requirements be directly related.

$$E_{\text{interfere}} = \frac{P_T K_{\text{path}} F_{\text{margin}}}{SINR \; R_{wc}^{\alpha_{wc}}} \tag{6-5}$$

The relationship of the receive link distance and minimum separation of nodes is:

$$S_{manual} = \left(\frac{R^{\alpha_{wc}} SINR}{F} \right)^{1/\alpha_{bc}} \tag{6-6}$$

A DSA system operating at the statistical typical case, has separation distance given by:

$$E_{\text{interfere}} = \frac{P_T K_{\text{path}}}{S_{DSA}^{\alpha_{tc}}} \tag{6-7}$$

Where:

S_{DSA} is the separation distance required for typical case, DSA spectrum deconfliction that meets the interference threshold

The relationship between the reliable communications range (R_{wc}) and the DSA and non-DSA separation distances is therefore:

$$S_{DSA} = \left(\frac{R^{\alpha_{wc}} SINR}{F_{\text{margin}}} \right)^{1/\alpha_{tc}} = S_{\text{manual}}^{\frac{\alpha_{bc}}{\alpha_{tc}}} \tag{6-8}$$

For a typical mobile radio, the best and worst paths are both diffractive for a path over five to ten km in any but the most extreme terrain. In practice, there is both a relative and absolute limit on the separation distance. For terrestrial communications, the transition to diffractive is typically around 1 km in outdoor environments.

The treatment of duty cycle is somewhat subtle. If the duty cycle is 25%, then does that mean that there is an opportunity to load 4 times as many radios? This is possible based on the expected value of spectrum usage, assuming: (1) that the system could use spectrum that was immediately available, (2) access can be deferred when spectrum is not available, and (3) a mechanism for effective use of temporarily excess spectrum is provided. This is certainly an appropriate model for a system that is intrinsically tolerant of access delays, such as a Delay Tolerant Network (DTN) [160], but for many applications would not be acceptable. The extent to which duty cycle (proportion of actual channel usage) can be exploited is a function of the size of the pool of spectrum. A pool of 10 channels shared among 40 users is quite different statistically from 100 channels shared among 400 users, in terms of the reliability it can deliver, even through the spectrum utilization ratio is equivalent. Therefore, two additional parameters are introduced to fully specify a shared spectrum environment; the required availability of spectrum (A_{Spectrum}), and the pool size (*Pool*). The availability of a specific number of channels (*needed*) at a given duty cycle and *pool* size (number of radios that share a pool of spectrum) is given by the binomial distribution. Each radio needing a channel is a trial, and a trial success is defined as a radio not needing a channel, therefore the probability of success is 1- *duty*. The availability is the probability that more than *k* radios out of *Npool* radios do not need channels, when the probability of needing one is *duty*.

$$A_{spectrum} = \sum_{k=needed}^{Npool} \begin{pmatrix} Npool \\ k \end{pmatrix} (1-duty)^k \, duty^{Npool-k}$$

$$\text{for } needed \leq N_{pool} \qquad (6\text{-}9)$$

Where :

A_{spectrum} is the probability of spectrum being available for any arbitrary request

For large pool sizes, this becomes a normal distribution, and a more convenient description of the CDF uses the regularized incomplete beta function (I_x):

$$A_{\text{spectrum}} = 1\text{-}P_x(X \leq needed - 1) = 1 - I_{1\text{-duty}}\ (Npool\text{-}needed\text{-}1\ needed) \qquad (6\text{-}10)$$

The benefits of a statistically large pool of spectrum are clear in the values of A_{Spectrum} for various values of relative spectrum availability, as shown in Figure 6-4, expressed as a ratio of mean demand.

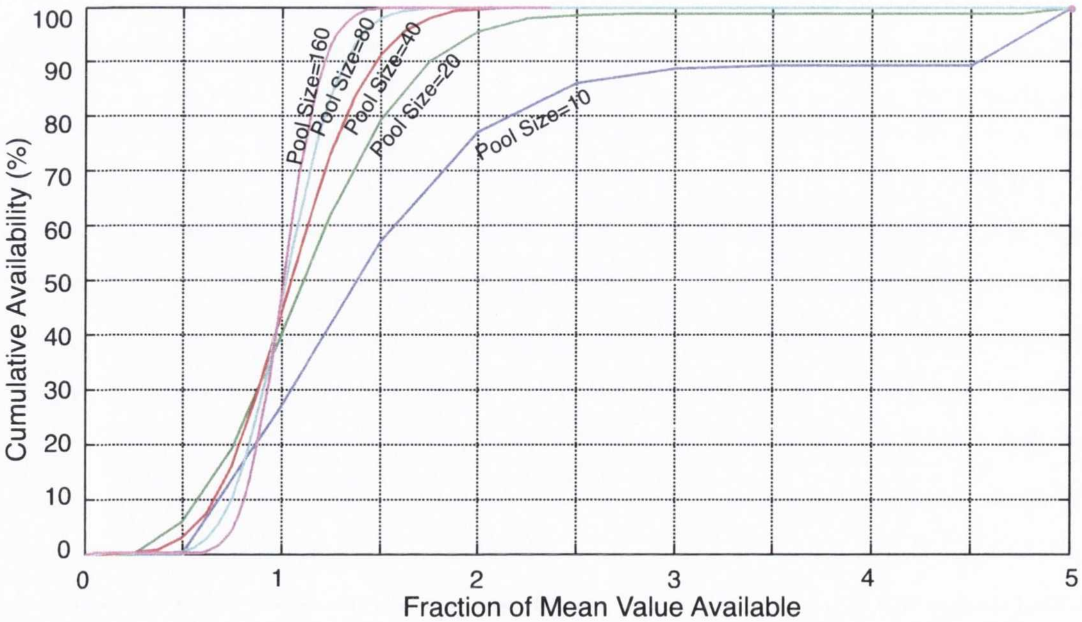

Figure 6-4 Representative Values of Spectrum Probability of Availability (A_0) for 20% Duty Cycle

Assuming that spectrum is provisioned linearly with the expected value of spectrum usage, without a sufficiently large number of radios sharing spectrum, the pool size is too small statistically to ensure access to spectrum at high enough confidence to support reliable operation. For example, in the above case of a 20% duty cycle device; high confidence operation ($A_0 > 98\%$) essentially requires dedicated spectrum for each radio in a pool of 10 radios, at least twice the mean value for a pool of 20 radios, but only 25% margin above the mean for a pool of 160 radios sharing the spectrum. Although the expected value of utilization (duty times units) creates an average spectrum usage, it is an inadequate expression of spectrum needs for smaller pools, when reliability is considered. To be effective, DSA requires statistically significant spectrum pools. All users benefit from large scale spectrum pooling. Note that spectrum pools do not need to be contiguous. In fact, the argument of Chapter 4 strongly favors spectrum pools that are highly decorrelated in respect to spectrum usage. As a minimum, the spectrum pool should be sufficiently varied to cover as much of the range of pre-selector settings as possible, to maximize the PSS value in Equation (4-14). An analysis of the issues associated with the transition to a pool-based approach to spectrum management for the Public Safety community, a potential early adopter, is provided by Le [161], Lehr and Jesuale [162] [163], Weiss and Jondral [164], and Jesuale and Eydt [165].

Another way of looking at the sharing of spectrum is to assume that there is an existing pool of radios present before the introduction of an additional set of DSA radios. The probability that a given pool shared with incumbent radios is sufficient to achieve opportunities for a DSA radio population is also given by Equation (6-9). The pool from which DSA frequencies are drawn need not be limited to DSA operation, but could be shared with non-DSA devices as well, increasing the density of spectrum usage, without implementing DSA in all of the devices that occupy the band, and increasing the pool size (and thus reliability) of the DSA devices.

The metric by which these will be assessed is *Spectral Density*. This metric is highly sensitive to individual designs, environments and usage patterns.

$$\text{Spectral Density} = \frac{\sum_{i=1}^{n} duty_i \cdot bandusage_i}{Bandwidth_0 \cdot Area_0} \tag{6-11}$$

Where:

n	Number of Users in the Spectrum
$duty_i$	The duty cycle of user i
$bandusage_i$	The instantaneous spectral bandwidth used by user i
$Bandwidth_0$	The total bandwidth made available to the set of users
$Area_0$	The geographic area over which the spectrum is used

Clearly, the optimum situation is to have spectrum available for use and allocation in all locations not within the interference radius of the potential receiver. This density is thus one device per interference area. The density of a de-conflicted mobile device is the operating area, plus the region surrounding the perimeter of the operating area out to the interference radius. Our measure is the ratio of the maximal possible device density over the achievable device density.

$$\begin{aligned} \text{Area Effectiveness} &= \frac{\text{Interference Area}}{\text{Deconflicted Area}} \\ &= 1 + \frac{2\pi r_{interference}^2}{Deconflict_{area} + r_{\text{int}\,erference}\ Deconflict_{perimeter}} \end{aligned} \tag{6-12}$$

Where:

$r_{\text{interference}}$	Interference Radius of the worst case link pair
$Deconflict_{\text{area}}$	The de-conflicted mobility area
$Deconflict_{\text{perimeter}}$	The de-conflicted mobility perimeter

These principles will be used to investigate several use cases in which spectrum selection will be studied, and in examining the effects of various dynamic spectrum algorithms. These will start with some unreasonable practices and situations, and then be refined to represent typical spectrum allocation practices.

The first case is a simple two-way communications between two vehicles that will navigate around the continental US. The communications range of the device is 7 km, so a reasonable estimate of its interference radius is 22 km. (Estimated by assuming propagation in an r^2 environment and a 20 dB SNR at the edge of the operating range.) Since the US has an approximate area of 9.6 million km^2, the Spectrum Effectiveness is 1.5 times 10^{-4}. It is even lower when considering the effect of duty cycle. If it is assumed that the vehicles operate for 8 hours a day, and transmit at most 10% of the time, the actual utilization drops to 5.2 times 10^{-6}! This is a strong reason to believe that the cognitive radio technology has an opportunity rich environment to improve these practices!

A more reasonable approach might be to only dedicate a segment of spectrum for a single city. Los Angeles has an area of 12,000 km^2. Adding the perimeter that must also be protected from use, the total reservation increases to 22,000 km^2. Using the same operating assumptions as above, the effectiveness is approximately 2.3 times 10^{-3}.

It should be noted that many systems of wireless devices already implement more advanced techniques than the baseline. For example, wireless hubs may search for open channels; cellular devices may have open slots or frequencies assigned; and 802.11a may adopt Dynamic Frequency Selection (DFS). These examples do not argue against this baseline; they represent the first (albeit simple) implementations of dynamic spectrum systems, and thus, cognitive radio. The cognitive radio clearly distinguishes itself from these more simplistic implementations when it fuses perceptions of all nodes impacted, controls these decisions by policy, and balances these decisions with the needs of the network which it supports!

6.4 Interference Tolerant DSA Operation

The previous section considered the initial case of DSA radios sharing spectrum with devices that were not tolerant of any interference to their operation, as they were presumed not to have DSA capability, and therefore, any energy in their communications channel was considered to be noise that would degrade their link margin. Such devices have no recourse when they detect the existence of interference. This section considers the case where the incumbent radios have DSA capability, and use that capability to not only locate open spectrum, but to mitigate the effects of interference through internal algorithms and behaviors. The effect of this change is profound: instead of having to create essentially near-zero interference operation, DSA is permitted to create a possibility of interference, with a probability constrained to be low enough so that the aggregate network costs of frequency relocation do not exceed the additional capability created by these more aggressive spectrum usage practices. This discussion is in contrast to the cognitive radio principles assumed by many researchers, such as outlined by Jovicic and Viswanath in their information theoretic analysis of cognitive radio, in which the secondary user could cause no interference to the primary user [166].

The approach of Spectrum Outage Probability (*SOP*) provides a model for determining probability that the aggregate signal strength from a set of homogenous emitters exceeds a set emissions mask at a fixed location within a network [167], [168]. Pinto shows that the *SOP* of the aggregate environment of an infinite plane of a Poisson distributed field of nodes is given by an Alpha-Stable distribution of the aggregate field [169]. Although this *SOP* formulation considers a range of frequencies and interference masks, only the intended transmission needs to be considered in this analysis, and for both simplicity and generality, a flat energy distribution within the transmission band is assumed, such as provided by an OFDM waveform. In addition, since only the decisions being made by a single node are being considered, this formulation can be simplified to the anticipated distribution of a single emitter.

The distinction of this analysis from the similar formulation of the Gupta and Kumar paper [141] is that there, the assumption was that an interference event caused a loss of throughput due to the failure of the information transfer. This analysis assumes that the consequence of a single or multiple interference events is a forced transition to a new frequency. The cost of this transition is not a failure to transfer information, but a temporary loss of capability while the transition is performed and the network is reestablished on the new (and presumably less interfering) frequency. The mean effect on throughput in a statistically independent environment is thus:

$$Capacity = \frac{T_{\text{sensing}}}{T_{\text{sensing}} + SOP \cdot T_{\text{rendezvous}}} \tag{6-13}$$

where:

SOP is the probability of outage due to interference
T_{sensing} is the interval between sensing intervals
$T_{\text{rendezvous}}$ is the time to re-rendezvous the physical layer, which includes the time to process the decision to move, coordinate the network on a new frequency, and re-establish any network state required prior to re-initiating throughput transfers

In fact, for short intervals of time (T_{sensing}), the assumption of independence is very conservative, as the movement of nodes is typically much slower than the sensing rate. This value is therefore an upper bound on the rate of disruption, and a lower bound on the capacity. A general treatment of $T_{\text{rendezvous}}$ is provided by Silva and Guerreiro [170]. This analysis is independent of the method of rendezvous applied, which can range from conventional control channels, even to swarm techniques [171], [172].

A unit-less generalization of this relationship can be constructed by substituting the relationship of the rendezvous time to the sensing interval, referred to as the DSA index.

$$Capacity = \frac{(1-SOP)}{(1+SOP)+(SOP \cdot I_{DSA})}, \text{where} \quad I_{DSA} = \frac{T_{\text{rendezvous}}}{T_{\text{sensing}}} \tag{6-14}$$

Figure 6-5 illustrates the relationship of capacity, *SOP* and the DSA index (I_{DSA}) for some representative values of each.

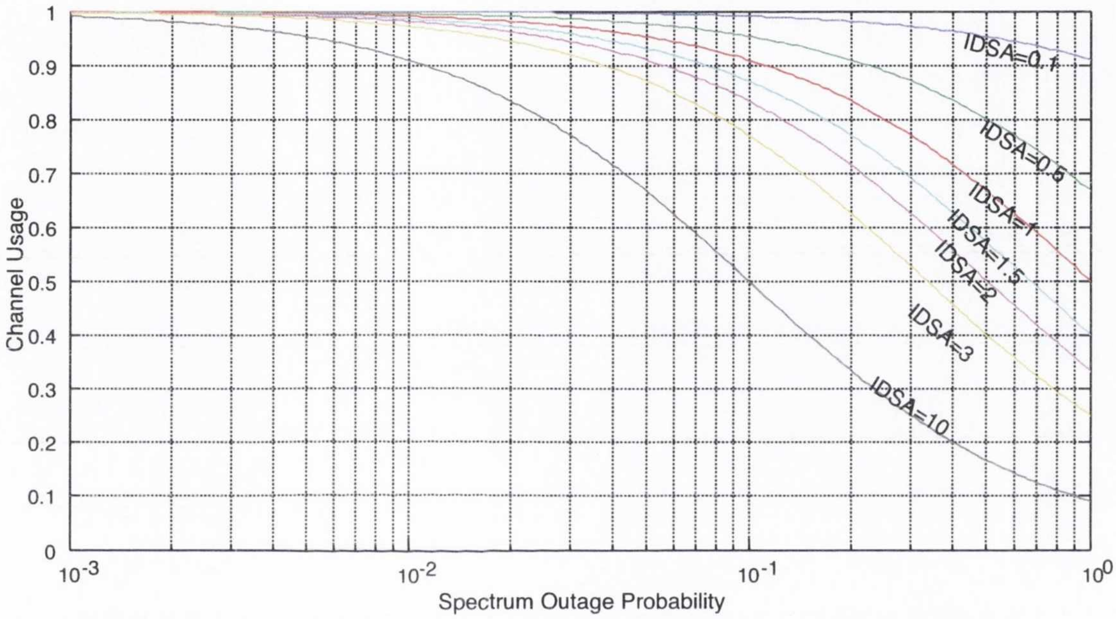

Figure 6-5 Capacity for a Typical Set of Sensing Rate and Rendezvous times

The figure demonstrates that frequent interference events can occur with minimal effect on network throughput so long as the networks can recover or relocate rapidly. As an example, the XG experiments are reported [71] to have a sensing interval of 100 ms and a re-rendezvous time of 186 ms, of an I_{DSA} of 1.86. This range of I_{DSA} values can operate at even a *SOP* of 10^{-1} with only a 15% degradation of channel access performance. While a 10^{-1} chance of causing interference would be unacceptable to a non-DSA radio sharing a band, it is not a significant performance impediment to a DSA radio sharing the same band. The opportunity provided by DSA on DSA spectrum sharing is to greatly increase the density of nodes and to maximize the aggregate throughput. The product of node density and channel access provides the total channel usage density per unit area and spectrum allotment:

$$D_{channel} = Density(SOP) \cdot \frac{1}{1+SOP \cdot I_{DSA}} \tag{6-15}$$

Where:

Density (*SOP*) is the node density for a given SOP value

In a generalized DSA mobile deployment (randomly distributed nodes), the probability that a frequency will receive interference at a given time is the probability that a node is in a given position, and that its path to the victim receiver has a propagation constant that is less than or equal to the one that would yield energy equal to the frequency abandonment threshold. Assuming:

c = channels available and randomly assigned

d = density (units/area x)

then, for a given Δx, the probability of at least one node being on a given frequency is:

$$P(\text{node on frequency}) = \frac{d}{c}\Delta x^2, \quad \text{for } \frac{d}{c}\Delta x^2 \ll 1 \tag{6-16}$$

Relocation only occurs if there is a path with propagation α below the value that would create a signal at the abandonment threshold at a receiver, referred to as α_{inrange}. Andersen et al. [155] shows that α is typically between $\alpha_{\min} = 2$ and $\alpha_{\max} = 5$. Any "required" $\alpha < 2$ indicates[42] that there is no chance of interference, $\alpha > 5$ means that there is high probability, or almost certainty, of interference, and between these values, that the chance is statistical based on α_{mean} and α_{variance}.

In this analysis, the propagation constant is considered the primary driver of interference level. The cited experiments also describe an additional multipath-induced fading term, but this is not considered here, since in situations of mobility, it is reasonable to assume that the nodes will be present in all possible states of the fading condition, and it is likely to be present in a non-faded condition for at least an instant of the period of time during which the emitter is operating.

The chance that interference is above the interference threshold is given by:

$$P(\mathit{interference}) = \frac{d}{c}\Delta x^2 \quad P(\alpha_{\mathit{inrange}} > \alpha) \tag{6-17}$$

Where:

$\mathcal{P}$ (α) is distributed between $\alpha_{\min}$ and $\alpha_{\max}$ with $\mathit{mean} = \alpha_{\text{mean}}$

Characterizing the distribution of α values by their mean and variance, and recognizing that there is a physical limit on the lowest value of α, but not one on the largest, then the probability distribution of α can be expressed as:

---

[42] This is for outdoor propagation where the alpha must be 2 or greater. Indoors, such an assumption may not be due to waveguide effects over shorter distances. In this case, the minimum value of alpha may be adjusted and the mean and variance of the alpha tailored to the specifics of the environment.

$$P(\alpha_{inrange} > \alpha) = 1 - \frac{\gamma\left(k_a, \frac{\alpha_{inrange} - \alpha_{\min}}{\theta_\alpha}\right)}{\Gamma(k_a)} \tag{6-18}$$

where:

$$\theta_\alpha = \frac{\alpha_{\text{variance}}}{\alpha_{\text{mean}}}, \qquad \text{k}_\alpha = \frac{\alpha_{mean}^2}{\alpha_{\text{variance}}}$$

To compute the maximal value of α that would lead to interference, the power at the receiver at range *r* is:

$$P_{\text{interference}} = \frac{P_T K_{path}}{r^{\alpha_{inrange}}} \tag{6-19}$$

Therefore:

$$\begin{aligned} r^{\alpha_{inrange}} &= \frac{P_T K_{path}}{E_{\text{interference}}}, \quad \text{therefore} \\ \alpha_{inrange} &= \log_r\left(\frac{P_T K_{path}}{E_{\text{interference}}}\right), \quad \text{and therefore} \\ \alpha_{inrange} &= \frac{1}{\ln(r)} \ln\left(\frac{P_T K_{path}}{E_{\text{interference}}}\right) \end{aligned} \tag{6-20}$$

The area of the circumference of a given *r* is approximated by 2πr Δ*r* for small Δ*r*. Therefore, the probability of Single Source Interference (*SSI*), that one emitter capable of causing interference above the threshold energy is:

$$P(SSI) = \frac{d}{c} 2\pi \int_0^\infty r^2 \left(1 - \frac{\gamma\left(k_\alpha, \frac{\frac{1}{\ln(r)} \ln\left(\frac{P_T K_{path}}{E_{\text{interference}}}\right) - \alpha_{\min}}{\theta_\alpha}\right)}{\Gamma(k_\alpha)}\right) dr \tag{6-21}$$

This equation is essentially the probability of a propagation exponent that is below the level that would cause interference. Closed-form integration of the Gamma function is not practical, but a numeric solution for this integral is straightforward. A typical base case with the link and environmental parameters shown in Table 6-7 is used in the following examples.

Table 6-7 $\mathcal{P}$(SSI) Base Case Parameters

Parameter	Base Case
K_{path}	-85 dB
$E_{interference}$	-105 dBm
P_T	1 dBm
α_{mean}	2.7
$\alpha_{variance}$	0.6
$\alpha_{minimum}$	2.0

The relationship of the node density and the probability of interference are shown in Figure 6-6 for a range of potential interference of one per 10^{-5} through 10^{-1} sensing events. As discussed previously, the units are a function of the degree of correlation and are highly situational. If all events were independent, this interval would be the sensing scan interval. If nodes are in motion, it would be the time to move between significant changes in propagation or changes in node usage.

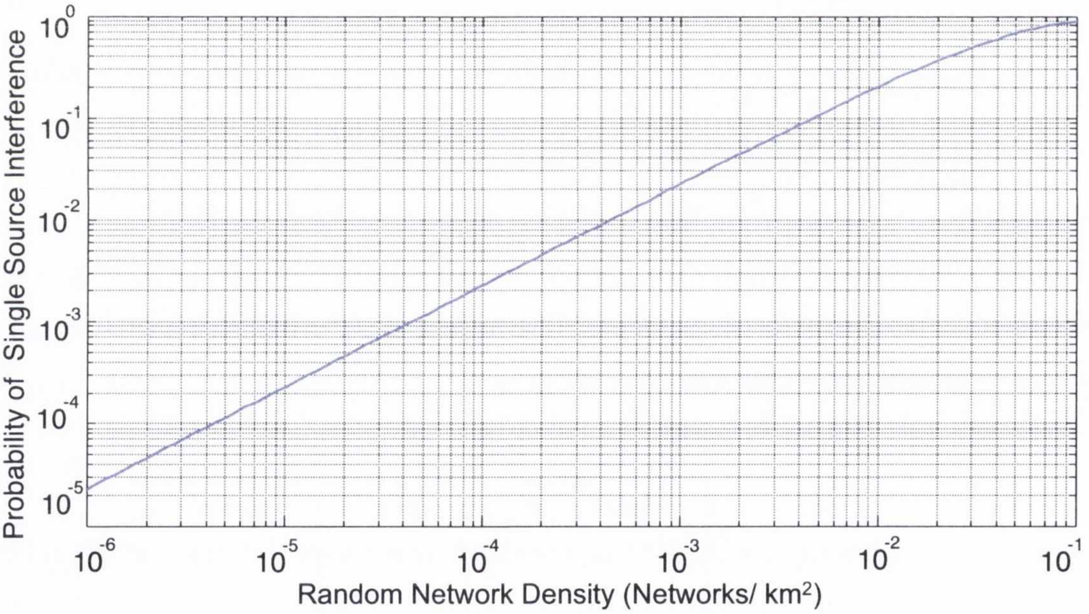

Figure 6-6 Relationship of Density and Probability of Single Source Interference of a Mobile Node

Assuming that nodes must relocate in the spectrum upon an interference event, a typical value of I_{DSA} of 1.75 yields the resulting individual node throughput characteristics shown in Figure 6-7.

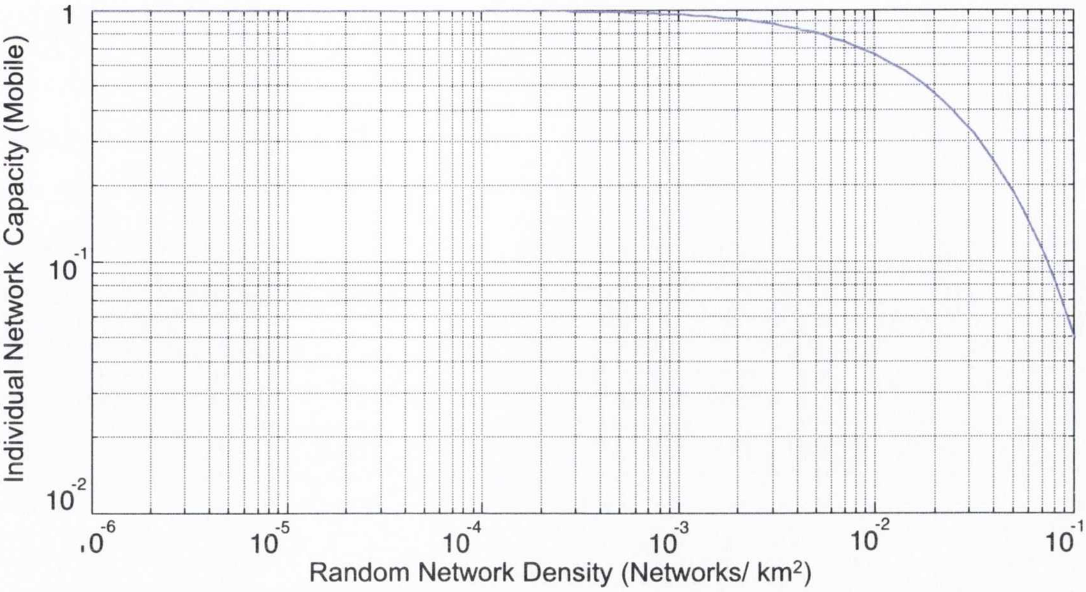

Figure 6-7 ***Mobile Node Throughput as a Function of Deployed Density***

Using the capability statistics from Equation (6-14), the aggregate capability for a range of densities is shown in Figure 6-8. This provides the relationship of node density, I_{DSA}, and the aggregate capacity over the region and spectrum available. The aggregate capacity (channel access per unit area and spectrum) can then be determined. For a fixed bits/hertz, this is in the desired units of throughput per unit area. The maxima of this function indicates that optimal aggregate throughput occurs at very high density, and thus very high interference probability, as show in Figure 6-6.

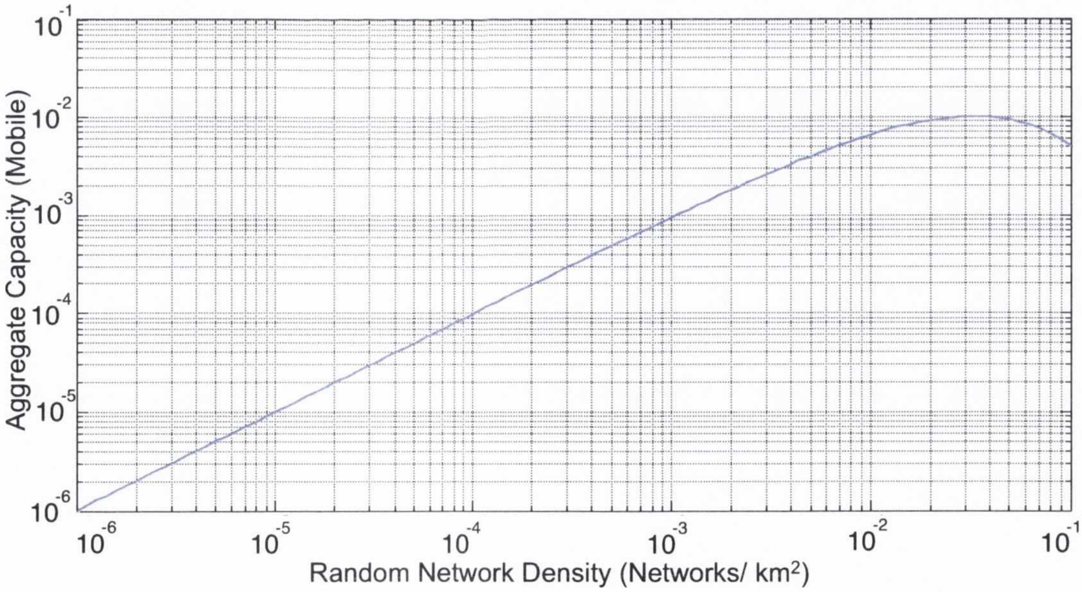

Figure 6-8 ***Relationship of Node Density and Aggregated Link Capability in a Mobile Environment***

The aggregate throughput benefits of adding nodes to the environment is much more significant than the loss of throughput due to spectrum relocation, up to the point where very dense environments are achieved, where an inflection point occurs, and the effect of additional nodes is to reverse this trend, and actually detract from aggregate throughput. This result (a maximum capability) is consistent with the theoretical analysis of Gupta and Kumar [141]. The SOP analysis shows that DSA is highly advantageous for reasons in addition to the spectrum management arguments typically put forth. It argues that it is important to not view DSA technology as intrinsically overlaid over non-adaptive radios (which forces the radios to be very conservative in avoiding interference) but to accelerate the development of dense cognitive radios in which the interference avoidance behavior is allocated to both sides of the spectrum sharing relationship, allowing much more aggressive spectrum sharing without impact on communications reliability. Although the specific benefits are case specific, in the example just discussed, the density of nodes can be increased by four orders of magnitude, and the aggregate carrying capacity of the spectrum increased by over three orders through dynamic management, rather than avoidance of interference.

A similar analysis can be conducted for stationary nodes, where a minimum distance can be enforced. In this case, the probability of interference is driven by the distribution of propagation exponents.

$$P(SSI) = (1 - \frac{\gamma\left(k_\alpha, \frac{\frac{1}{\ln(r)}\ln\left(\frac{P_T K_{path}}{E_{\text{interference}}}\right) - \alpha_{\min}}{\theta_\alpha}\right)}{\Gamma(k_\alpha)} \tag{6-22}$$

In this situation, the driving characteristic is the variation in the propagation environment between the interfering and victim node. The probability of interference as a function of node spacing for a fixed (or enforced minimum separation distance) arrangement of two homogeneous nodes (as described previously) is shown in Figure 6-9.

There is a fairly rapid drop off of interference beyond a threshold distance. In the mobile environment, the chance of interference was treated as independent; i.e., the chance was the same for every sensing cycle. This is extremely conservative, as mobility patterns are typically significantly slower than the scan rate proposed for DSA operation. In the case of mobility, there are very few causative agents for rapid changes in the mean value of the channel. In this case, the effect of the I_{DSA} index is much less significant than the degree of correlation of the channel propagation conditions. "Effective I_{DSA}" can be adapted to consider the correlation of the channel.

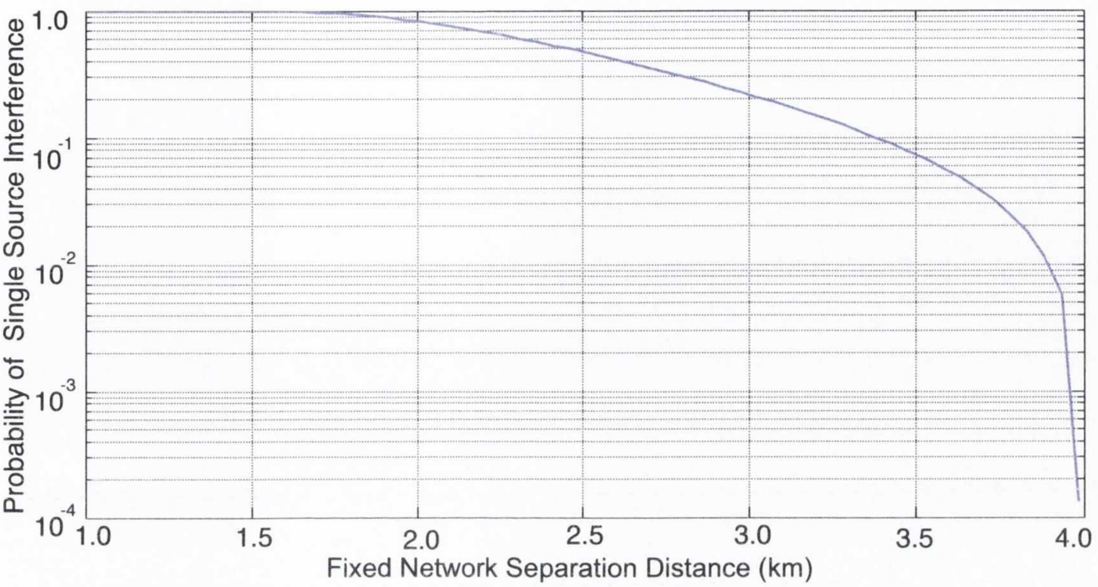

Figure 6-9 Probability of Interference in the Base Case for Fixed Nodes

Figure 6-10 includes very low values of I_{DSA}, reflecting this high degree of channel time correlation.

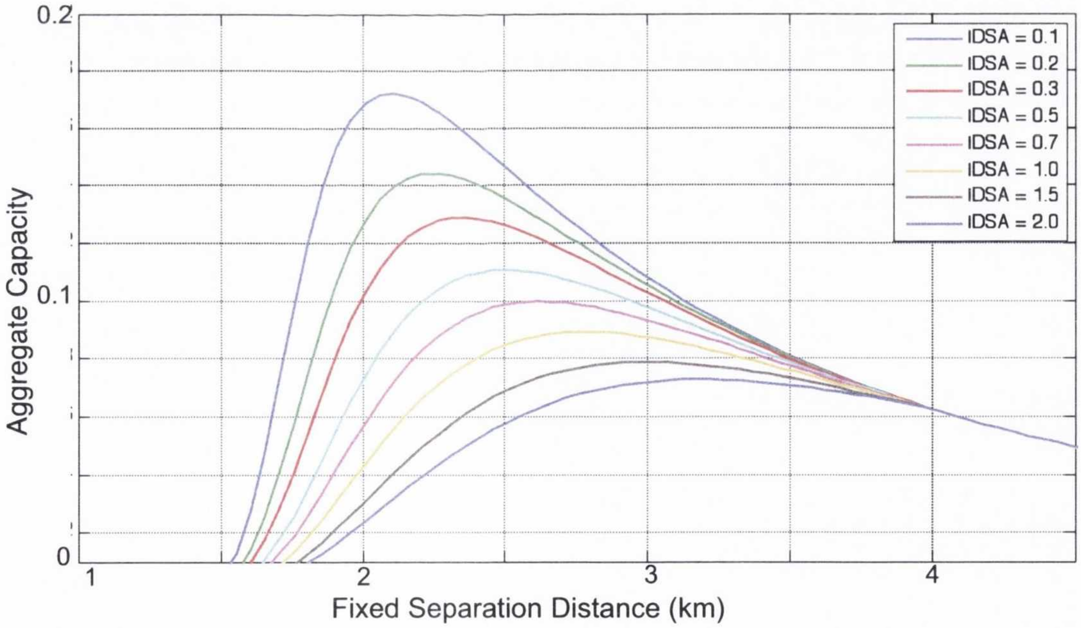

Figure 6-10 Aggregate Capacity for Fixed Nodes in the Base Case with a Range of IDSA Values

This chart illustrates the same effects as shown for mobile nodes. Very rapid spectral relocation of network operations enables significant capacity increases (over twice the density of nodes) to be achieved through adaptive interference mitigation, rather than interference avoidance. Selection of an appropriate value is thus highly dependent on the local causes of channel variations. An effective I_{DSA} of 0.1 can be considered to be a spectrum usage condition (considering mobility, channel characteristics, etc.) that changes at one tenth the speed at which the DSA system can re-rendezvous on a new frequency.

The optimal density of nodes is significantly dependent on the channel correlation. In this base case, when the channel is slow compared to the DSA operation (I_{DSA} in the 0.1 to 0.5 range), the density of nodes can be increased from the range of one 4 km spacing, to 2.1 km spacing, an increase in node deployment density of over 3.5 times.

6.5 Dynamic Spectrum Access Benefits

The preceding sections provided the quantitative basis for assessing the usage density relationships of DSA and non-DSA operating modes. The benefits in wireless device and information density, and typical benefits, are summarized in Table 6-8.

Table 6-8 Quantitative Benefit Categories for Dynamic Spectrum Access

Benefit	Usage Density	Typical Density Benefit
Static Nodes	Density Increased by: $\left(Area_{manual}\right)^{\frac{2\alpha bc}{\alpha tc}}$	Density increased proportional to manually planned exclusion range (in km)
Mobile Nodes	Density Increased by Equation (6-21)	Typically 3 orders of magnitude
Reduced Duty Cycle	Dynamic spectrum usage approaches a reduction of $DutyCycle^{-1}$ with sufficiently large pool of spectrum, and for sufficiently long periods of non-use	Increases density of devices by *DutyCycle* for sufficiently large sets of nodes

6.6 Dynamic Spectrum Access Conclusions

The use of DSA has been shown to be both essential and beneficial to the implementation of adaptive and cognitive radio systems. The benefits of DSA extend far beyond just better utilization of fallow spectrum, but can be employed to address some of the limiting factors that constrain wireless network capability.

Specific conclusions of this chapter include:

1. There are quantifiable benefits from DSA in terms of reducing separation between users of common spectrum. These benefits relate directly to risk management allowances such as fade margin and best and worst case propagation assumptions. There is an irreducible minimum separation required to assure that interference is below the required SNR of the protected signal.

2. Significant pools of spectrum are required to assure high confidence, real time access to spectrum through DSA mechanisms. The expected value (*radios times dutycycle times bandwidth*) of spectrum is only approachable with statistically significant pools of spectrum.
3. The complexity of multilateral spectrum selection is not of significant benefit if operation is constrained by threshold energy based spectrum policies. However, if there are no occupancy policies, the perception of devices within link distances may be quite different, and therefore these mechanisms are essential to operation without imposed thresholds.
4. Although DSA has significant benefits in sharing spectrum with non-cognitive devices, its greatest benefits may accrue from the ability to mitigate interference, thereby reducing the necessity for additional "risk" margin, since any interference that does occur is mitigated by one radio relocating. The effectiveness of this approach is a function of the rate of change of the channel and node locations compared with the rendezvous time of the DSA mechanism. However, exploitation of this approach will require coordination mechanisms throughout the communications architecture, and, optimally, the design of applications that are supported by the network services.

Chapter 7 Spectral Footprint Management

7.1 Spectral Footprint Management Objectives

This chapter addresses the selection of waveform and operating mode options in order to minimize the use of spectral resources that could otherwise be used by other users of the spectrum, again based on the assumption of the flexibility provided by DSA operation. This chapter considers spectrum resources as more than strictly the occupied bandwidth, but also considers the geographic extent over which use of that spectrum is denied to other users. ***The critical issue in spectrum optimization is not spectrum use, it is spectrum reuse***.

This recognizes the opportunity cost of an emission to other users. A spectrum resource is considered to be the product of the occupied bandwidth and the area over which the interfering signal is above a threshold of interference. While a single user of the spectrum might be indifferent to the effect of waveform choices on other users, the users or operator of a network of nodes might be more concerned with maximizing aggregate capacity of the nodes, rather than the individual ink data rates.

This is consistent with the increasing complexity of wireless systems. Initially, spectrum was dedicated to a single transmitter or link pair, so was a scalar of bandwidth. With the advent of TDMA and CSMA MAC layers, the spectrum utilization became a product of time and bandwidth. With the advent of cognitive radio and DSA, it is now more appropriate to think of spectrum usage as a volume; with dimensions of time, bandwidth and denied area. The proper units of spectrum cost should be considered to be seconds-Hertz-kilometers^2. This metric is fundamental to the design of complex networks.

There are two environmental aspects to the selection of frequencies for cognitive radio operation. The well known one is fitting into the spectral footprint of other, non-cooperative radios, which is the subject of much of the current research. The second, and more subtle one, is to minimize the radio's own footprint. For years, modulation developers have defined spectrum efficiency as the number of bits per Hertz of bandwidth, and have often used this metric as a scalar measure of the best modulation. The assumption has been that the design that utilized the least spectrum was intrinsically less consumptive of precious spectrum resources. Presumably, the more bits per Hertz, the more spectrally efficient the modulation. This proposition is worth examining.

The problem is how to define spectral efficiency metrics. Classically, digital radio engineers have attempted to minimize the spectrum used by signals through maximizing bits per Hertz. This is a simple and readily measured criterion. We broaden the view of spectrum impact to include not only the amount of spectrum used, but also the area over which it propagates until it is no longer a significant interference contributor in comparison to the noise floor. It is worth noting that volume would be a more generalized measure, but since most spectrum usage is on the surface of the earth, consideration will be limited to two, rather than three dimensions. Shannon [173] shows a basic relationship between minimum energy and maximum possible bits per Hertz. The Shannon bound argues that essentially infinite bits per Hertz can be achieved, but it can only be accomplished by increasing the energy per bit exponentially, and thus, the radiated spectral power increases at the second power of the spectral information density.

An immediate observation is that as the bits/Hertz increases, the energy required per bit goes up exponentially. The proportionate cost in energy of going from one bit to two bits per Hertz is essentially the same as that required going from six to seven. Theoretically, essentially arbitrary bits per Hertz are possible if the channel is sufficiently clean, and power is available. But, it is tough on battery-powered devices! It is equally tough on the other users of the spectrum. In order to increase the bits per Hertz, the radio must now transmit both slightly more bits, and vastly more energy per bit, in order to meet the E_b/N_o requirements of the receiver. The relationship between the area within the communications region (where E_b/N_0 is achieved) and the area within the interference region (where the interference threshold is reached) is important. Figure 7-1 illustrates the strength of a signal at the boundary of the communications region, and at the boundary of the interference region. This example depicts an interference threshold where Signal Energy = Noise Energy, or 3 dB signal plus noise over noise.

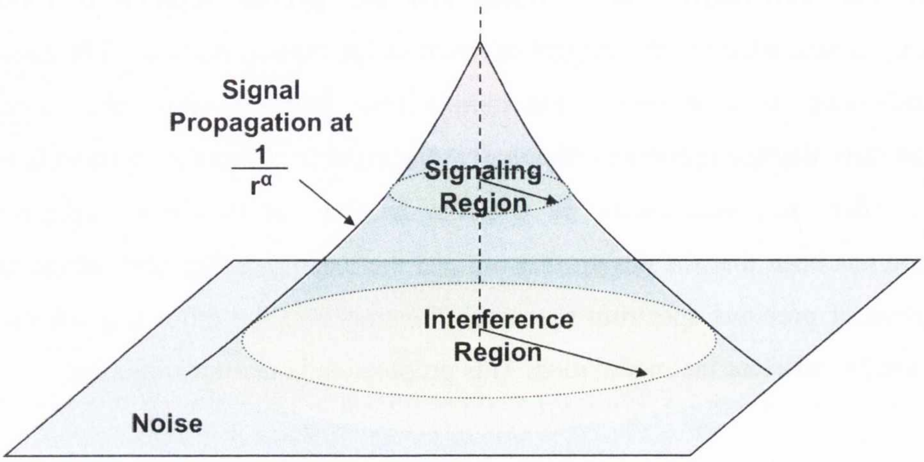

Figure 7-1 Relationship between Communications and Interference Regions

Using the Shannon channel capacity equation [173], the increase in spectral energy required to increase spectrum efficiency can be computed. If simple spectrum were the measure, such a radio strategy would be effective. However, for most frequencies, spectrum reuse is a valuable characteristic when normalized by the area of the region over which its use by other devices is precluded. Therefore, "***bits per unit area***" is an appropriate measure for how effectively spectrum is used. Extending this metric to consider propagation is important to understanding how increasing bit density denies spectrum reuse to other radios. Ignoring channel effects, such as multi-path and absorption, propagation between terrestrial antennas can be (simplistically) modeled as r^{α}, with α typically varying from 2 to 4, depending on distance, antenna height, and frequency [174].

Figure 7-2 illustrates the impact of modulation order on likely footprint for a range of modulation orders and propagation conditions. Even small increases in spectral efficiency have major and disproportionate impacts on the footprint of radio transmissions.

Figure 7-2 Impact of Bits per Hertz on Spectrum Footprint

Clearly, simplistic spectrum strategies that address only bits/Hertz are not effective, and are quite counterproductive for spectrum users as a whole. The better the propagation conditions (i.e., the lower the value of alpha), the more counterproductive it is to increase the number of bits per Hertz from a spectrum optimization viewpoint. For most propagation conditions, it is clear that by increasing the spectral efficiency of one radio, it disproportionately reduces the ability of the spectrum to support other users. Since alpha generally increases with frequency (for propagation in VHF and above), the implication is that more sophisticated strategies are needed. Increasing modulation constellation depth is a poor solution for this radio, as it greatly increases the required energy. It is an equally poor solution for the other radios sharing the spectrum, as it essentially raises the noise floor throughout an increased region, and therefore, precludes operation at a rapidly increasing radius from the transmitter. The effect of modulation order on the carrying capacity of a system of densely deployed nodes is quite severe.

This leads to a measure that recognizes this inherent compromise, or trade space. In this case, bits per Hertz/area reflects that the critical issue in spectrum optimization is not spectrum use; it is spectrum reuse. Measurement and optimization should not focus on how the radios perform individually; but on how the radios allow other radios to share the spectrum in a close to globally optimal manner. The relative value of this measure is shown in Figure 7-3, for a range of simple propagation conditions. The Interference Area represents the region over which the interference exceeds the threshold, and thus, the opportunity cost. Bits/Interference Area is the value of the throughput obtained, divided the by opportunity lost. It is normalized to the cost of a one bit/Hertz link.

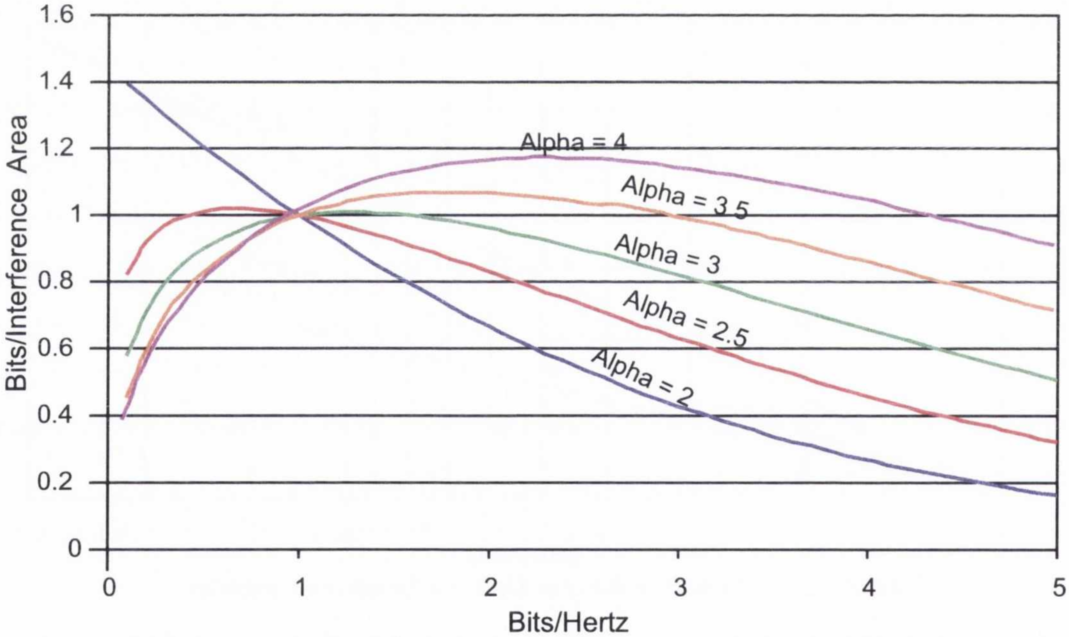

Figure 7-3 Impact of Constellation Order on Bits per Hertz/Area

This graph compares the product of the bits transmitted times the relative area over which interference is present, compared to the area over which it is present using one Bit/Hertz, Shannon-optimal waveforms. At 2.5 bits/Hertz, a node would have an increase in bits/Area effectiveness of 20% if the propagation exponent (alpha) was 4, but would have a loss of bits/Area effectiveness of 50% if the propagation constant were 2. A wireless node in an r^2 environment is only 50% as efficient in going from 2 to 4 bits per Hertz, whereas, in an r^4 environment, the same mode change is efficiency neutral. This adds another dimension to the algorithms for a cognitive radio. If it, or its spectrum access protocol or etiquette, is attempting to optimize overall spectrum utilization, then awareness of the relative antenna height and range is required in order to determine between alternatives that are either significantly beneficial or dramatically detrimental! Measuring propagation loss to just the intended receiver is not optimal, since the critical issue is how fast the signal is attenuated on the far side of the receiver, and in other directions.

Logically, if there is a 12 dB SNR, then the signal is above the noise for another four times the link range in the r^2 environment, and only one range distance when the propagation is r^4. The preceding discussion assumed that the power was perfectly matched to the channel; in practice, such an alignment is impossible to achieve or maintain over any degree of time varying channel, so real interference levels will fall above these values.

To address the transition from a radio focused optimizing strategy to one that considers the effects across all users of the spectrum, it is necessary to introduce a waveform selection strategy that is "***α-aware***". This strategy considers the propagation of the radiated energy beyond the communications region of the radio, until it is so dissipated that it has no effect on other devices. ***While conventional power management addresses strategies between the transmitter and the receiver, α-aware waveform selection is cognizant of the effect beyond the receiver out to all other devices in the spectrum. In effect, this construction introduces the "opportunity cost" consideration in the spectrum optimization process, by awareness of the propagation physics.***

When discussing footprint management, it must be recognized that the application of this optimization is applicable to those situations where the radio derives some benefit from being cooperative to the other users of the spectrum. While many authors have assumed that shared spectrum access would resemble the historic "*Tragedy of the commons*" [175]-[178], there are fundamental reasons to believe that even uncooperative users would not benefit from intentionally or unintentionally destroying the utility of the spectrum for all other users. One dramatic difference is that radios stop "polluting" once the transmission stops. There is not the permanent degradation that was implicit in the derivation of the original term. In fact, there are two situations where it is reasonable that a radio would desire to control its footprint. The first is that the radio is part of a larger regional network, and there is a high probability that the network would benefit from decreased footprints. The other is when there is an established or mandated protocol or etiquette that requires or causes the devices to minimize their impact on other users. Work by Neel also is supportive of the conclusion that destructive operation is not inherent in the behavior of self-interested devices in the spectrum [16]. Collectively, all radios benefit if each is provided adequate spectrum to use low-order modulations, even at the expense of locally increased spectrum bandwidth usage.

7.2 Spectral Footprint Reference Evaluation Metrics

Simplistic models of the interference ranges of the cases to be considered are shown in Figure 7-4. The use of high spectral efficiency (high energy per bit and thus high energy per Hertz) waveforms reduces the number of nodes within a given space to accommodate the increase interference radius.

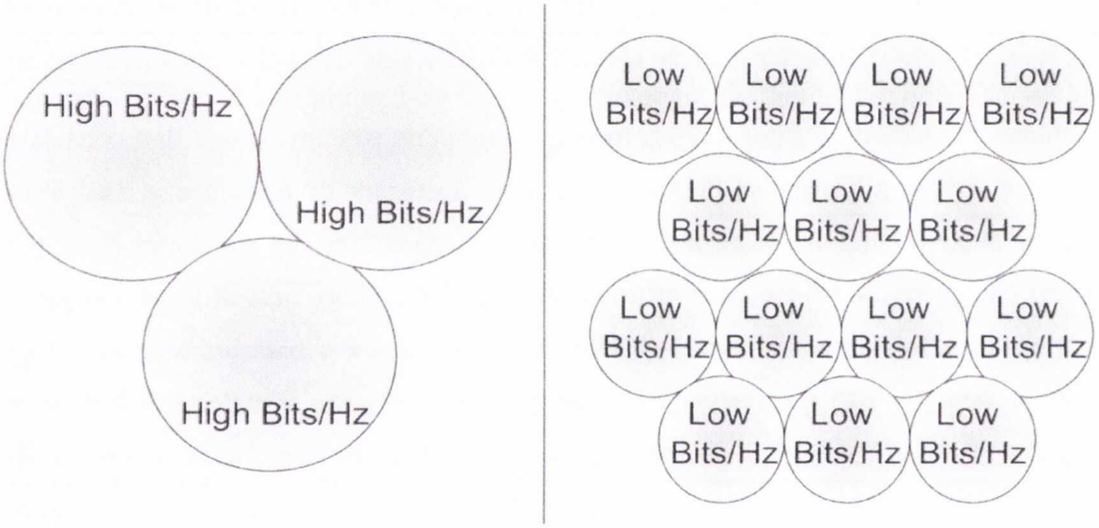

Figure 7-4 Spectral Efficiency and Spectral Footprint

The optimization of spectral footprint is not a scalar consideration. It must balance several factors, including the effect on other users, and the costs to the device. Taking the theoretical performance of a Shannon-limited receiver, known propagation exponents, and constraints of noise threshold, the bound on spectral/area effectiveness can be determined from Shannon as follows [173].

$$C = B\log_2\left(1+\frac{S}{N_0}\right) \tag{7-1}$$

Where :

C = Channel Capacity
B = Channel Bandwidth
S = Signal Energy Level
N_0 = Noise Energy Level

This is approximated by:

$$C = \log_2(10)\, B \log_{10}\left(\frac{S}{N_0}\right), \quad \text{for } S/N > 1 \tag{7-2}$$

$$\text{Therefore,} \quad \frac{S}{N_0} = 2^{\frac{C}{B}} - 1 \tag{7-3}$$

For narrowband waveforms, such as typical practice for mid and long range communications devices, the interference radius (in this case for $S=N$) can be considered to be proportional to the α^{th} root of the difference in power.

$$r \propto \left(\frac{S}{N}\right)^{\frac{1}{\alpha}} = \left(2^{\frac{C}{B}} - 1\right)^{\frac{1}{\alpha}} \tag{7-4}$$

Where :

α = Propagation exponent, as in r^{α}; A = 2 for free space, typically around 3.8 for close to the ground, diffracted environments [155]

r = Range of communications link

The Spectral Footprint (area) ratio of an idealized Shannon-bound waveform to the equivalent 1 Bit/Hertz can be expressed as:

$$\text{Spectral Footprint} = A \propto \left(\left(2^{\frac{C}{B}} - 1\right)^{\frac{1}{\alpha}}\right)^2 \tag{7-5}$$

The metric under consideration is the information per area that is transmitted compared to the region over which operation of other devices is precluded or detrimentally impacted. Normalizing Equation (7-5) by the spectral bandwidth (B) yields the area effectiveness (Spectral Information Effectiveness (SIE)) of a bandwidth strategy.

$$SIE = \frac{C}{AB} = \frac{C}{B}\left(\left(2^{\frac{C}{B}} - 1\right)^{\frac{1}{\alpha}}\right)^{-2} \tag{7-6}$$

Where:

SIE = Spectral Information Effectiveness, which has units of bits per Hertz per unit area

The optimal operating point is a function of propagation exponent. To determine the optimal value of bits/Hertz as a function of the propagation exponent, the SIE term is differentiated with respect to the modulation rate (C/B) yielding:

$$\frac{\partial\ SIE}{\partial\left(\frac{C}{B}\right)} = \frac{1}{\left(-1+2^{\frac{C}{B}}\right)^2} - \frac{2^{\left(1+\frac{C}{B}\right)}\frac{C}{B}Log(2)}{\left(-1+2^{\frac{C}{B}}\right)^3} \tag{7-7}$$

The root of the derivative of the *SIE* ratio (at other than the degenerate endpoints) yields the optimal operating point. The optimal modulation bits/Hertz for a given alpha is reflected below:

$$\frac{C}{B} = \frac{\alpha + 2\ \text{ProductLog}\left(-\frac{1}{2}\alpha\ \text{e}^{\frac{-\alpha}{2}}\right)}{2\ \text{Log}(2)} \tag{7-8}$$

Specific values for a range of propagation exponents are shown in Figure 7-5. Note that this chart matches the extrema for the numerically derived curve in Figure 7-3. This relationship is one that is immediately usable in cognitive radio.

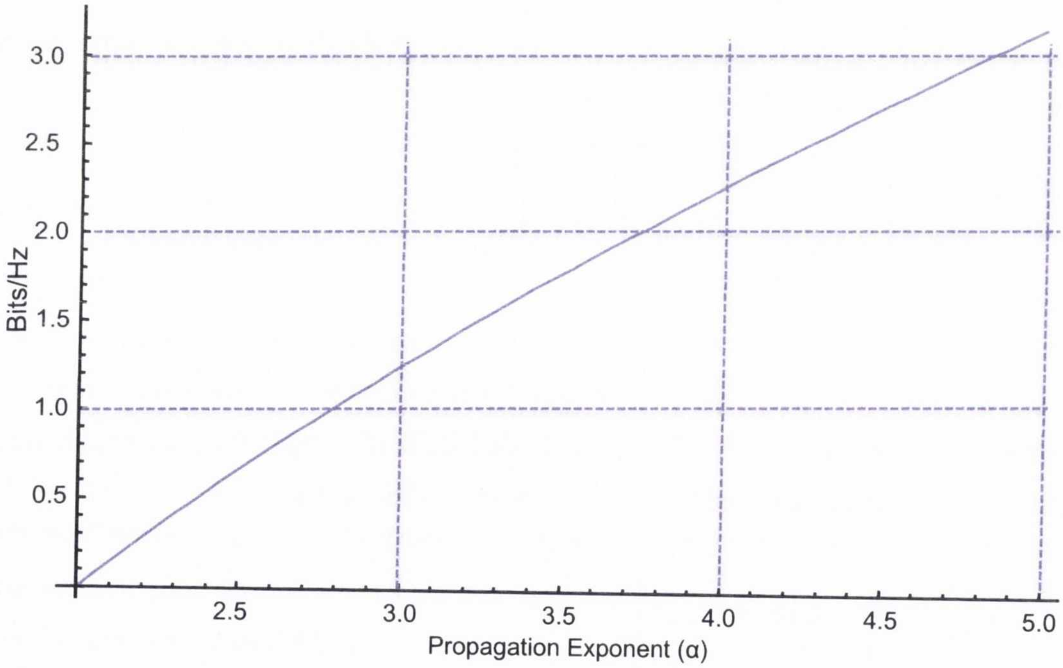

Figure 7-5 ***Optimal Values of Bits per Hertz for a Range of Propagation Values***

7.3 α-Aware Waveform Selection Algorithms and Methods

Optimization of spectrum footprint is a straightforward process if the path characteristics are known or imputable. Imputing them is readily accomplished if the transmitting power, receive power and distance between the nodes is known. Otherwise, the terms must be assumed based on the likely employment of the nodes. Of course the node cannot be aware of the exact propagation to other users of the same spectrum, but it can be aware of its own placement (high or low antenna, attenuating environment, or not), line of sight operating concept, etc., which provides a reasonable basis for imputing a value of α. Once an estimate of α is available, the footprint optimization can be readily accomplished in accordance with Equation (7-8). ***The effect of these strategies is to transition from spectrum use as the focus, to spectrum re-use as the focus and criteria for optimization.***

7.4 α-Aware Spectral Footprint Management Benefits

The benefits of this technique are essentially unbounded, since it would be possible for a device to elect to operate at essentially arbitrary modulation order. In practice, the range of modulations typically used have spectrum footprint differences of almost 1,000, with only a corresponding increase in bit rate of 10. Benefits in spectrum information density from several times to almost 100 are achievable through the management of this footprint.

7.5 DSA and Spectral Footprint Management Impacts on Network Scaling

To consider how a cognitive network might make use of the flexibility of DSA and cognitive adaptations, it is worthwhile examining the fundamental issues of scaling raised by the work reported by Gupta and Kumar [141]. In this model of a MANET, the nodes are routed across a large (multiple hop) space in which interference arises with contended use by one or two hop neighbors, since the two hop neighbors have sufficient signal to raise the noise floor, even if not sufficient for reliable demodulation. This model is shown in Figure 7-6.

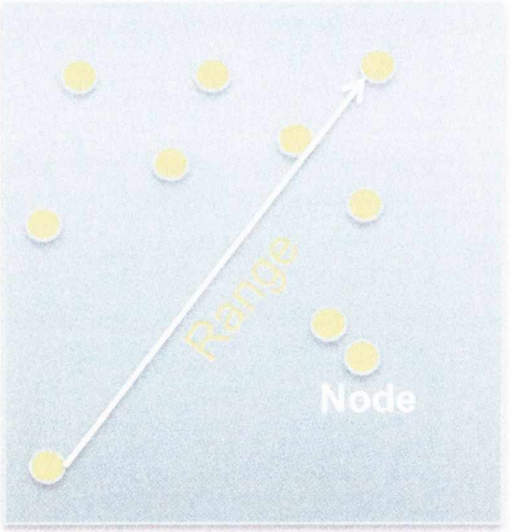

Figure 7-6 Conventional View of MANET Operation

This design has the fundamental design issues pointed out by Gupta and Kumar. For reliable communications, only one node can transmit at any one point in time, hidden nodes cause jamming that MAC layer protocols cannot resolve, and the capacity of an individual node is constrained to $n^{-0.5}$, at best.

These inherent constraints of fixed operating modes can be somewhat addressed by adaptive mechanisms. Adaptive networks with DSA can dynamically manage networking within the same region by partitioning the nodes by frequency, with a high probability of reducing interference from two hop neighbors by ensuring that DSA algorithms only rendezvous nodes into clusters with total reachability between all of the members of the cluster. Therefore, there are no "hidden" nodes that can cause interference, but yet are not able to reliably participate in all MAC layer exchanges.

A simplistic example of this structure is shown in Figure 7-7, where the adjacent sub-networks all have unique frequencies. This approach initially divides up the spectrum pool, but; as will be shown, for large networks, this provides more stable scaling characteristics. For simplicity, this is depicted and analyzed with the assumption that interference range is approximately twice the communications range; a reasonable assumption, but one that will be removed later in the analysis. Subnets are of a fixed membership size, so as more nodes are added to a given region, the link range is made shorter, and more sub-networks are configured. Network links only need to communicate within the individual block, and the network can simultaneously allow transmissions in all blocks, so the aggregate capability is that of sixteen networks, each sharing one-quarter of the original spectrum, without the need for complex, overhead inducing, and ineffective MAC-layer coordination. The problem with this organization is that the independent networks are just that; independent and not connected into a network fabric.

Figure** 7-7 **Partitioned MANET Sub-networks

To resolve this constraint it is necessary that each node provide at least two receivers that are operable independently. An additional set of frequencies is overlaid on this network, so that its connectivity is "half a phase" from the original lay-down, so that the connectivity created connects adjacent sub-networks, as shown in Figure 7-8. This structure now requires eight frequencies, but has thirty-two independent networks operating, so the net effect is four times the throughput.

As the space or node count to be networked increases, the number of frequencies required to network remains fixed, but the number of independent networks increase, so there is an inherent scaling in this architecture. An intelligent network strategy would not have to establish arbitrary standards of network size or separation, and could vary these as traffic and latency constraints dictated.

The first case to consider is if spectrum is unconstrained, and the network operates to maintain a constant spectral power density at the receiver. In the following discussion, the strategy is that the number of sub-networks that are operated is allowed to vary, but the number of nodes per sub-network is held approximately constant.

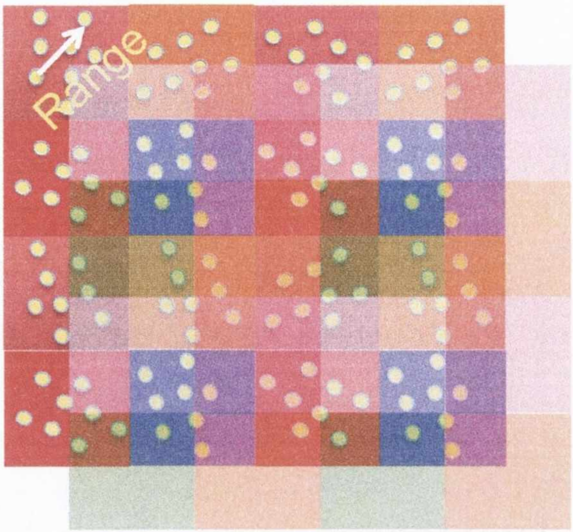

Figure 7-8 Interconnected MANET Sub-networks

Individual throughout (*I*) is proportional to the total network capacity (*N*) divided by the number of nodes (*n*) divided by the number of hops (*h*) needed to deliver messages.

$$I \propto \frac{1}{n}\frac{1}{h}N \qquad (7\text{-}9)$$

Total network capacity (*N*) is proportional to the number of networks, which are proportional to the square of the range *(r*) on a surface, and the information rate of the links I.

$$N \propto \frac{1}{r^2}R \qquad (7\text{-}10)$$

Since the Rate of the links in unconstrained spectrum is proportional to energy to be distributed across the spectrum evenly, follows the square (or higher order in diffracted regimes α_{Comm} law with distance), the link rate I is directly proportional to the range *r*.

$$R \propto \frac{1}{r^2} \qquad (7\text{-}11)$$

The range (r) to the average closest neighbor is proportional to the square root of the number of nodes, therefore:

$$r \propto \frac{1}{\sqrt{n}} \qquad (7\text{-}12)$$

The Hops (*h*) are proportional to the range of the individual links *I*, since as nodes are added, packets must transverse through more nodes to reach the same point.

$$h \propto \frac{1}{r} \propto \sqrt{n} \tag{7-13}$$

For closely spaced nodes ($\alpha_{Comm} = 2$), the throughput for a single node ideally scales with the square of the number of nodes, so long as spectrum resources are available to the network. For diffracted paths, the capacity increases more rapidly with *n*.

$$I \propto n^{\frac{(\alpha_{Comm}-1)}{2}} \tag{7-14}$$

This limit is obviously not reasonable as *n* approaches infinity, since infinite spectrum is not available, and practical processing on the devices has a finite limit. Never the less, this is meaningful, since it demonstrates that as network size increases, the capability is significantly increased, which sets a much higher baseline from which the spectrum constrained performance (described next) case departs.

The second case occurs when available spectrum is exhausted, and the extra signal power available must be applied to increase modulation constellation depth, since the energy is contained within a fixed spectrum. The throughput in this case is constrained because the increase in sub-network information rate is constrained by Shannon's Limit [173] to grow logarithmically with energy, rather than linearly. Additionally, increased spectrum power density would lead to an increase in the distance between networks required to avoid inference against a fixed noise energy to bandwidth criteria.

For a constant EIRP, the energy (*S*) available to the receiver is the inverse of the range to the α_{Comm} power.

$$S \propto \frac{1}{r^{\alpha_{Comm}}} \propto r^{-\alpha_{Comm}} \tag{7-15}$$

According to Shannon, the Rate (*R*) of the link with constant bandwidth is proportional to the log of the energy.

$$R \propto \log_2(S) \propto \log_2\left(n^{\frac{\alpha_{Comm}}{2}}\right) \tag{7-16}$$

The spacing to adjacent networks for constant interference energy, is a function of radius I and the spectral power density (SPD), which is proportional to *S*, and propagation constant for interference distances, which may be different than for the communications distances.

$$spacing = r + dr$$

$$\frac{dr}{r} \propto S^{\frac{1}{\alpha_{Int}}} = r^{\frac{-\alpha_{Comm}}{\alpha_{Int}}} \qquad (7\text{-}17)$$

The number of networks is proportional to the square of the distance between nodes.

$$nets \propto \left(\frac{1}{r\left(1+d\right)} \right)^2 = \left(\frac{1}{r + r^{\left(1-\frac{\alpha_{Comm}}{\alpha_{Int}}\right)}} \right)^2 \qquad (7\text{-}18)$$

Substituting the relationship of nodes and link range provides the relationship of the number of nets and nodes.

$$nets \propto \frac{n}{\left(1 + n^{\frac{\alpha_{Comm}}{2\alpha_{Int}}} + n^{\frac{\alpha_{Comm}}{\alpha_{Int}}} \right)} \qquad (7\text{-}19)$$

The Individual node throughput I, the network capacity (*N*) divided by the number of sub-networks, times the throughput I, divided by the average transit distance; the square root of *n*.

$$I \propto \frac{n}{\left(1 + n^{\frac{\alpha_{Comm}}{2\alpha_{Int}}} + n^{\frac{\alpha_{Comm}}{\alpha_{Int}}} \right)} \frac{1}{\sqrt{n}} \log_2\left(n^{\frac{\alpha_{Comm}}{2}} \right) \qquad (7\text{-}20)$$

Two interesting cases are when propagation to both intended and interfered with nodes is identical, and when it is unbalanced. When the two alpha terms are in the free space region, and therefore similar; the capacity scaling relationship (for large *n*) becomes:

$$I \propto \frac{\log_2 n}{\sqrt{n}} \qquad \text{(for large n)} \qquad (7\text{-}21)$$

This is better than the non-dynamic alternatives by the Log term. When the propagation to distant nodes is greater (for example, the closest link partner nodes are free space propagation, and the distant nodes are in diffracted propagation, then the ($\alpha_{Comm}/\alpha_{Int}$) term is one half the free space value, and the scaling relationship (again, for large *n*) is $\log_2(n)$. Significantly, greater propagation losses to potential recipients of interference significantly change the rules of network scaling, and provide a significant benefit from the use of DSA in these situations.

Yet a third option is to use power control, and reduce power proportionate to range in order to maintain the same SPD at the receiver. This avoids increases in the exclusion region (*dr from equation 7-17*), but constrains the network to the same communications rate I. This is Equation 7-20, except setting the *R* term to 1, reflecting no additional link throughput or bandwidth usage.

$$I \propto \frac{1}{n}\frac{1}{\sqrt{n}}\frac{1}{r^2} \propto \frac{1}{\sqrt{n}} \tag{7-22}$$

In practice, it would be expected that different parts of the network would be operating in different regions, due to differences in spectrum availability, density, traffic loading, and propagation conditions. Therefore, likely performance probably falls in between these two extremes.

Figure 7-9 illustrates the relationship of the scaling effects for:

- Unconstrained spectrum
- Constrained spectrum with propagation constants of 2 (equal α)
- Bounding, and Power Managed Operation (1/SQRT(n))
- Generic 1/n curve
- Propagation to link partner equal to the potential victim (Equal alpha)
- Significant differential in propagation to link partner and victims (Near/Far)

With significant path loss differential (Near/Far line), the scaling performance of the MANET approaches the spectrum-unconstrained case, as the effects of interference drops off rapidly past the effective range of the radio.

Measurements reported by Andersen et al. [155] show a mean value for σ of 2.7; but with high variability, even within a small region. This points to the fundamental advantage of cognitive radio; that it can adapt and utilize different techniques in different situations within a single network. Where propagation is difficult, the network can compensate by making more use of high order modulation, where propagation is unimpeded, it can adapt power to manage the interference footprint. In practice, propagation from low antennas follows the fourth power rule after only several hundred meters, so the Near/Far curve is probably most indicative of typical scaling in the spectrum constrained scenario.

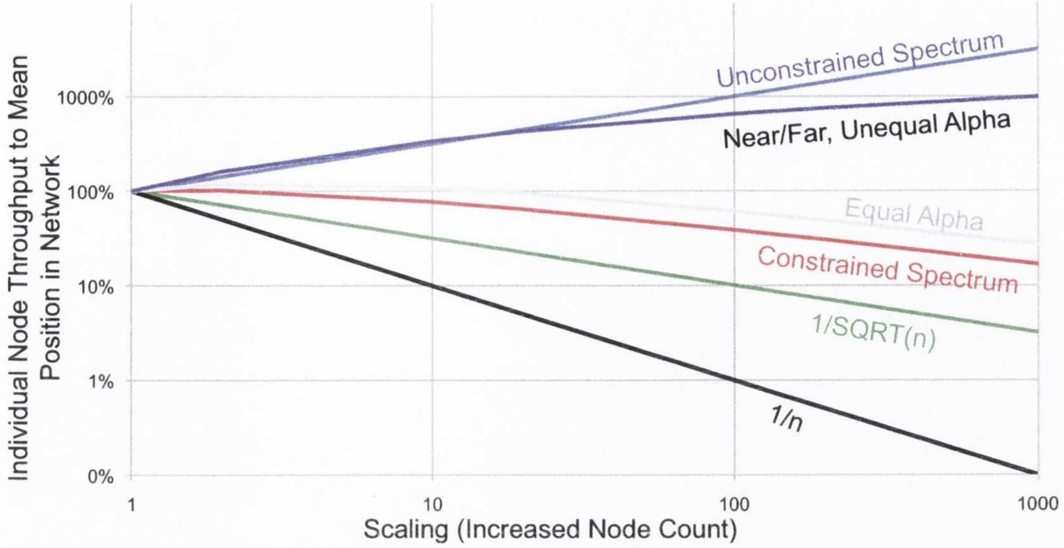

Figure 7-9 ***Illustrative Performance Curves for Various MANET Scaling Relationships***

7.6 Spectral Footprint Management Conclusions

The assumed optimizing strategy of maximizing bits per Hertz is demonstrably ineffective in maximizing the usage of the spectrum in a shared environment. In fact, in most environments, it leads to counterproductive levels of exponentially increasing energy and corresponding interference regions. There is a fundamental difference between optimizing the spectral effectiveness of a single node and optimizing a collection of nodes. The optimal spectral utilization of a single node is unbounded, while that of a network of nodes is shown to be discrete, with strong sensitivity to the propagation environment.

A relatively simple set of environmental awareness parameters can enable the cognitive radio to determine operating points that optimize the carrying capacity of the spectrum. Operating modes that minimize power, even at the cost of locally increasing spectrum usage have positive benefit for all nodes within the shared environment. There is a fundamental partitioning of techniques that optimize the individual link usage of spectrum and those that optimize the global use of spectrum.

Implementation of this approach is a fundamental shift in how the wireless community considers spectral efficiency. This requires a transition from a radio or link-centric analysis to one that considers all of the occupants of the spectrum as a single system.

Chapter 8 Extension of Principles to Network Level Decision Making

8.1 Implications on Network Level Decision Making

The advent of DSA has significant consequences on the organization of the network and its services. This chapter will investigate network adjustments that are uniquely, or significantly enabled by the dynamic topology and bandwidth organization possible through the spectrum and bandwidth adaptations discussed previously.

The preceding chapters have provided a quantitative argument that lower layer cognitive radio adaptations have a positive and profound effect on link performance, and/or affordability. These lower layer adaptations also provide a corresponding opportunity for profound changes in the upper network layers (as has been suggested by the author among others [107], [179] -[181]) that could enable network technology that would itself further enhance capability. These additional methods can be applied in concert with, or at a layer above the unilateral physical and media access decision making previously discussed.

Despite the inherently situational nature of these benefits, the possibility of these techniques provides a strong argument that likely cognitive radio performance can significantly exceed the bounds developed in the previous chapters by logically extending the framework described in Chapter 2 to upper layers. The methods described here are transitional, as research emphasis evolves from the physical layer focused cognitive radio, to the more expansive and inclusive cognitive Wireless Network Device, with a focus on the network and upper layers. This device has not yet been fully described in the research literature, but is the logical next step in the technology path.

This chapter will investigate the benefits to the device and to the network of violating the principles of layer abstraction. The simplest case assumes that a cognitive radio has insight into at least the transport layer identity of the upper layer traffic offered to the device. A more complex case is possible when the node also has knowledge of some aspects of the topology by which packets are being delivered, and can consider these facts in determining the operating point of the device[43].

---

[43] Note that these assumptions are not universally available, such as within an encrypted pipe, such as provided by Internet Protocol Security (IPSec) within a Virtual Private Network (VPN) when the transport layer features would not be accessible to intermediate systems, such as routers.

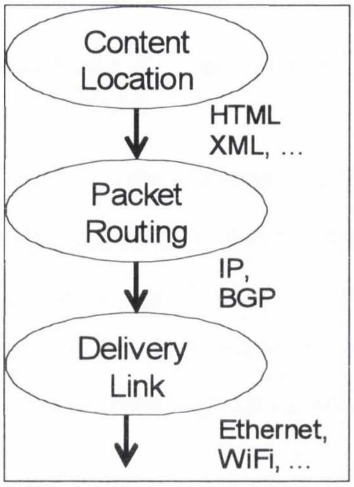

Figure 8-1 Illustrative ISO-Based Network Layer Interactions Model

Figure 8-1 illustrates the classical, ISO-Layer centric view of the management processes of modern networks. They essentially comprise three abstracted layers: one managing connectivity; one managing networking; and one managing content persistence and delivery.

They are abstracted by well understood and standards-based interfaces, such as Ethernet and Wi-Fi, Open Shortest Path First (OSPF) and Transport Control Protocol (TCP) [182], and HTML and XML. Each of these domains is a separate management domain, and in most cases, are independent providers and equipment owners.

The advantage of this architecture is that an XML server does not need to be aware that content is being delivered over an Ethernet, Wi-Fi or other physical connection. Although there might be benefits in such awareness, the inherent scalability of the fixed infrastructure makes such awareness unnecessary. It can be argued that much of the success of the Internet derived from the abstraction of these layers, so that the link, network, transport, and application layer technology could evolve in parallel.

In wireless networks, there is no such inherent scalability. Spectrum is finite, and Shannon shows it is not scalable with finite or constrained energy. Interference effects correlate the behavior of large numbers of nodes. Protocols, such as TCP/IP, are highly effective in managing transport in the core Internet, but are much less effective when they cannot distinguish between congestion, limited bandwidth, high packet loss rate, and link disruption. Similarly, when content is widely accessed from statistically independent users, there is little argument to move the content around the network dynamically, but when content is generated and consumed by users whose locations are highly clustered, then it implies that the network could benefit from locating such content as close to the user cluster as possible, if it could be made aware of the source of the information needs.

The opportunity provided by cognitive radio is to make each domain of the information exchange process aware of the constraints and opportunities within other layers of the network. This peer-to-peer and collaborative architecture is depicted in Figure 8-2. In this architecture, no single domain makes unilateral decisions regarding the best way of organizing itself; benefitting from inputs from the other domains. For example, routing is

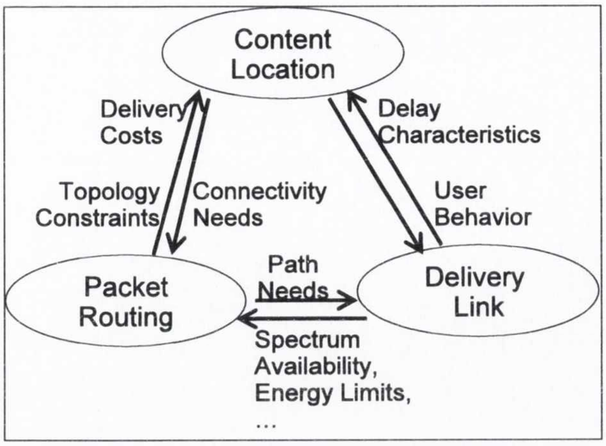

Figure 8-2 Representative Management Domain Interactions for Cognitive Networks

not determined solely through a network component, but is made aware of spectrum constraints for planning possible topologies; the correlation of users of the wireless network is provided to application layer services so they can locate content more closely to the likely users (example; maintain content on devices closest to the area of map coverage) and information flow characteristics are provided to the spectrum planning component to search for more appropriate spectrum, all simultaneously maximizing information delivery within the constraints of the equipment and the environment.

The physical layer techniques provided in Chapters 4 through 7 make this operational concept possible. The resulting device is not just a cognitive radio, but also a cognitive network information node. The following sections discuss specific techniques that are enabled by awareness and flexibility in the physical layer, and are appropriate to the dense networks that can evolve from affordable wireless networking hardware.

8.2 Dynamic Bandwidth Topology

Chapter 5 discussed the opportunities for Dynamic Spectrum assignment in a cognitive radio. This chapter extends that concept to dynamically vary the bandwidth associated with the communications link, presumably in response to instantaneous demands on the network. The benefit of this is that within a given spectrum usage, the dynamics of the user demand can be reflected in the bandwidth allocated (and de-allocated) to each link. Although the core networks of the Internet are planned with the principle that network demand is Gaussian or Poison, as one moves closer to the edge of these networks, it is intuitive that the traffic demand becomes less Gaussian; particularly, highly time correlated. A cell phone pattern of use is highly correlated; if it is in use at one second, it is highly likely that it will be in use the next second, and if not in use; it is highly likely to remain not in use in the next second. Even packet services have this characteristic, with users going through alternating page retrieval and reading phases, and completely inactive modes. With thousands of users sharing portions of the Internet, the sum of these variations are Gaussian, but at the wireless edge, they are far from independent. It has been demonstrated that the assumption of independent traffic is not valid for many classes of Internet traffic (which is much more self-similar than independent), and this situation is certainly much more likely in wireless links, due to the inherent relationship of location and information interests [183].

Leveraging this characteristic has been difficult because most wireless systems have a fixed spectrum environment, forcing the network to be constructed of fixed and invariant pipes. Much of the wireless networking infrastructure is derived from the technologies first developed for, and employed in, the fixed Internet. Therefore, the ability to adapt bandwidth dynamically is not present in these wireless systems. Even the wireless-specific technology developed for Mobile Ad Hoc Networking (MANET) focuses on routing decisions and awareness, not topology optimization.

This next section considers how cognitive radio wireless network nodes can balance spectrum and throughput bandwidth to respond to the dynamics of the traffic. This is one capability that wireless networks uniquely enjoy and can exploit. The simplest version of this technology would be a wireless node or relay point that varied the spectrum in use as the traffic queue in the router or a client terminal increased, and relinquished that spectrum to other access links when lower bandwidth usage could meet the Quality of Service (QoS) thresholds and minimize queuing.

A key question in examining link bandwidth adaptation is the time constant of the network traffic. Fixed wired networks typically re-plan their bandwidth on long-term intervals, which reflect static characteristics of the traffic demand. Exchanging planning, execution, and routing information has both time (delay) and resource (cost) implications, so the cognitive radio must constantly balance the temporal correlation of the traffic load, the cost of adaptation, and benefits of adaptation. As in dynamic spectrum, the benefits of dynamic bandwidth allocation can be utilized to either improve capability through matching bandwidth to traffic (with fixed signal constellations), or in lowered total energy usage through varied constellation sizes due to matching spectrum availability to bandwidth needs. Work by Sutton et al. demonstrates that bandwidth decisions can be communicated to receiving nodes implicitly by transmitting nodes without the necessity to discretely communicate it in advance, and can be integrated into the DSA rendezvous process [184], [185].

8.3 Cognitive Radio Enabled Dynamic Topology and Network Organization

While much of the argument for cognitive radio has focused on the characteristics of physical layer adaptation; it can be argued that the network layer, leveraging the flexibility enabled by the DSA and cognitive radio model, provides equally compelling benefits. DSA, by eliminating the specific assignment of individual links to specific frequencies and bandwidths enables adaptation not only of the physical layer, but also in the organization of the network itself. WNaN envisions the adaptations shown in Figure 8-3 to address and mitigate the effects of a range of environments, device limitations, and best satisfy information delivery requirements [6].

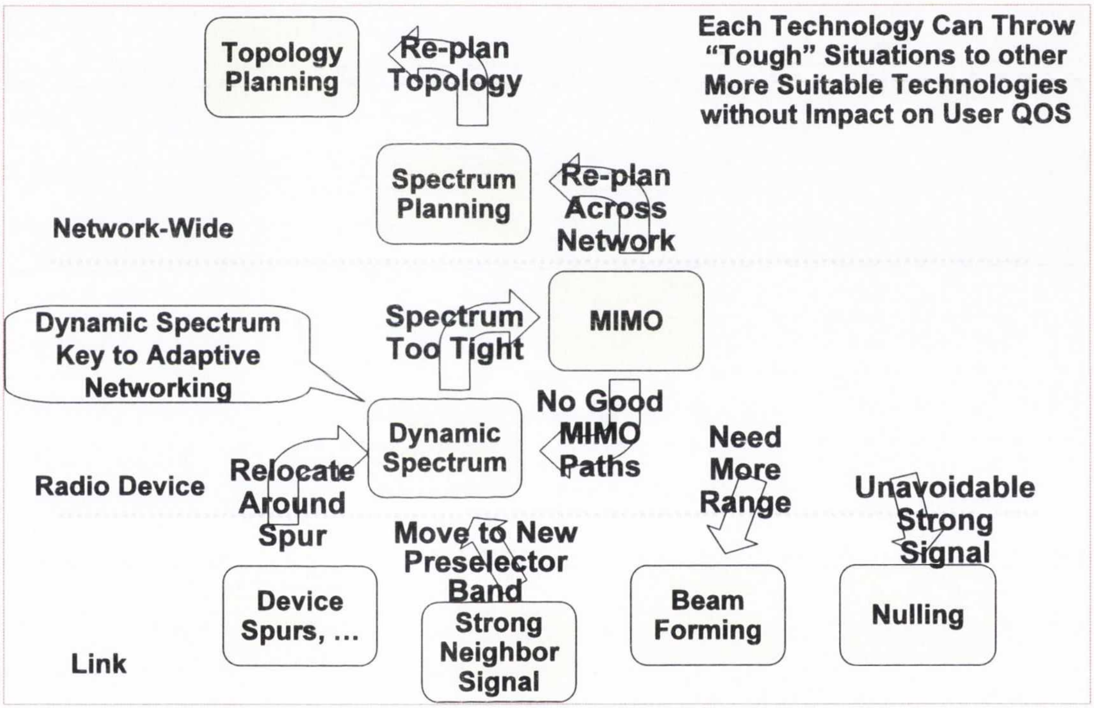

Figure 8-3 Upper Layer Exploitation of DSA-Flexibility

These adaptations have a prerequisite that the network be able to autonomously select frequency and bandwidth for each link. If bandwidth in one region of the network is inadequate, the network can "move" spectrum through locating or reassigning edge spectrum to supplement the throughput. If a link is interfered with by other signals or overloaded by adjacent channel energy, the cognitive radio locates spectrum more suitable to its operation, and is free to implement that decision locally and unilaterally. If there is no way to locate enough spectrum in a congested region of the network, the network changes its topology to route out of, and around, the congested region.

The DSA community has advocated the acceptance of DSA technology to address many of the shortfalls in manual frequency planning and management. This approach additionally argues to adopt DSA for the flexibility it affords:

- to the device for self-management of the environment, and the resulting reduction in component performance requirements, and thus cost;
- to the spectrum manager, to not have to separate receivers from strong sources of adjacent band energy; and
- to the network and applications, to reconfigure and re-provision wireless bandwidth dynamically.

8.4 Cognitive Radio Enabled Content-Based Networking

A fundamental premise of current practice has been to abstract the network and applications layers so that their technology, behavior and operation is independent and isolated, just as in the network layers themselves. Certainly, this abstraction has been one of the keys to growth of technologies such as the worldwide web, which have grown asynchronously with the underlying networking technologies. However, wireless is an environment where bandwidth is scarce and costly in terms of energy, spectrum, and equipment. It does not readily scale in capacity by incremental resources.

Disruption Tolerant Networking is an application of the DTN Research Group (DTNRG) architectural framework. The architectural framework of a DTN is provided by Cerf [160], and some experimental deployments in high performance wireless systems have been described [11], [186], [187]. The objectives are not only to mature the capabilities of DTNs, and their applications in networking, but also to use them as a network information metaphor, in which access to information is provided based on descriptions of content, rather than descriptions of the end nodes on which it is located [104]. In fact, this method of Internet access is quite typical, as Internet users resolve sources of information through search engines, rather than URLs. The proposed framework makes such access intrinsic, rather than a centralized application layer overlay.

There are several reasons for believing that this construct has advantages for a cognitive wireless network. In a fixed network, it is reasonable to assume complete advanced knowledge of the topology, the bandwidth available at nodes, and the characteristics of access needs. After all, fixed network connectivity does not suddenly appear; it takes discrete action and planning to make changes in the network. In a wireless network, there is little that is known in advance. In fact, the network may not be able to maintain full awareness of its own state at all times and places. Decisions on content positioning within the network are best made as information become available, and should be revisited as conditions change.

Introduction of DSA significantly enhances the networks ability to implement demand responsive topologies based on traffic and spectrum conditions, and to proactively implement topology changes through changes in link characteristics. It follows that the organization of information on such a network should be dynamic and interactive with organization of the network, spectrum availability, and user behavior. Equally, the organization of the network should be driven by the location of the content. Instead of putting one "organizing principle" above the other, a wireless network can leverage the interaction of these considerations to continually adjust the network, based on the applications, content, and the network's perception or understanding of the physical world's constraint on physical organization (spectrum, power, energy, interference, etc.).

Interaction with DTN's is performed using a bundle metaphor. A bundle is an integral grouping of metadata and data. Unlike a packet, it provides a context for any node that processes it. As a first step in the process of developing a content-based network, DTN research is experimenting with the use of caches to maintain opportunistic copies of content at all nodes that route, process, or receive it. They retain content as it passes through them. In the case of wireless nodes, the bundle may not even be intended for the node; it might have been addressed to another network member, but is retained if the transmission is overheard [87], [104]. There are significant benefits to this approach, particularly if user behavior has correlation. Demand for certain types of information, such as maps, local attractions, phone calls, directory services all would appear to be correlated. A similar approach demonstrated major increases in reliability by hop to hop caching in the Active Reliable Multicast (ARM) approach, which caches content at nodes in the multi-cast tree to avoid retransmit requests flooding up the multi-cast tree [188].

An additional benefit of this approach is that it allows the introduction of a mode referred to as "Late Binding". In late binding, a request for, or provision of, information is not required to be resolved to a specific network or node address. Instead, the request or content is "launched' towards the general direction of the ultimate destination, and the specific destination is determined as the bundle approaches the destination region, or as content is encountered and recognized as the request approaches likely repositories. This approach has the advantage that details regarding name spaces, subnet structure and reachability need not be propagated through the network, or be available at the time the information is generated or requested. Nodes and networks that are disconnected can address content and have content directed to them, and local names and other meta-data need only be distributed locally within a DTN region, rather than globally. This is similar to postal mail, where the full extent of the local delivery method are not needed or known at the initiation of the mailing. Another viewpoint is that late binding is a generalization of the single hop namespace resolution provided by email services.

An extension of the late binding is Content-Based Access (CBA) in which the information content of bundles is provided in separate metadata descriptions that are routed and cached along with bundles. The description of nodes, content, or content interests (searches) is symmetric and follows identical syntactic and semantic constructs. For example, a multicast group could be described as "*all nodes within 1 km of <lat, long>*", while a request for content could be "*the <name> of all nodes within 1 km of <lat, long>*".

Much network thinking assumes that bandwidth is hierarchical, with high-speed in the core, and lower speed at the edge, such as depicted in Figure 8-4 (a). A more realistic model is that the edge is often rich in internal connectivity, as shown in Figure 8-4 (b). Enterprise systems utilize this latter model by positioning high bandwidth services, servers, and file stores within the enterprise edge. For wireless systems to also leverage this local bandwidth, the model of a wireless network must shift from a collection of clients of Internet core-resident applications and storage, to self-contained, and self-reliant networks, with considerable inherent capability for distributed content management.

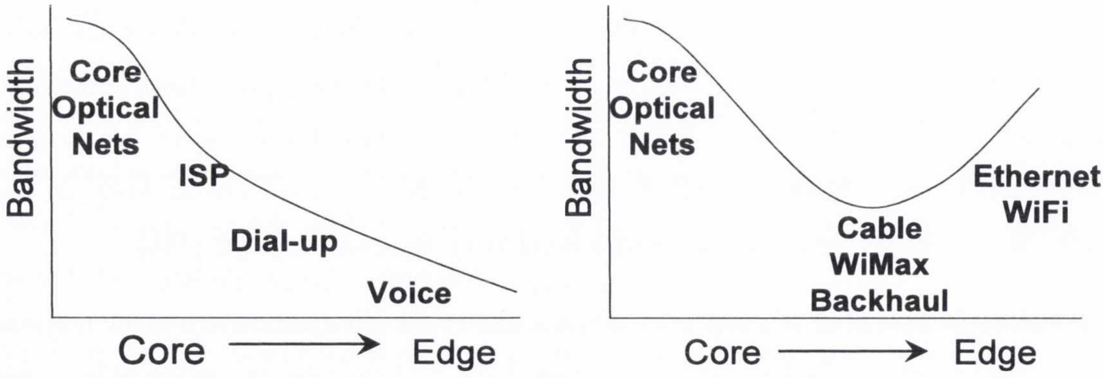

Figure 8-4 Alternative views of core and edge (a) Assumed, (b) Reality

A future possibility is that intra-edge bandwidth will often be plentiful and that the performance constraint will not be in the edge itself, but in linking the connected edge to the connected core [87]. Distributed, peer content caching is one technology to reduce the edges' dependence of the core, reduce the backhaul bandwidth needs, and evolve to peer-to-peer architectures that are rich in range of content. Proposed changes to the bundle specification will enable separate encryption for the bundle as a whole; the metadata, and payload data to provide confidentiality [189]. Nodes can provide caching or server-type services for content they cannot actually access, providing a "black side" cache or server. Notably, this opportunity is maximally exploitable if the network can adjust its organization to reflect changing effectiveness of the caching process by expanding and contracting the spectrum usage and waveform characteristics as needed.

DTN caching and late binding were recently assessed in a variety of wireless networking scenarios. The ability to cache content, defer address resolution, and maintain connectivity in disrupted environments was shown to improve wireless performance by factors of five to ten in the target wireless systems [186]. The concepts incorporated in DTN have been individually applied in systems for many years; what is unique is the integration of these features into a single pervasive network service.

In initiating the DARPA Wireless Network after Next WNaN cognitive radio program, the author postulated eleven theses for the program; as shown in Table 8-1 [87]. Although many of these theses relate to the network layer operation; they are enabled by the DSA and topology adaptation.

Table 8-1 WNaN Cognitive Radio Theses

1	The Network Will Adapt to the Mission and Organize Itself Responsively to Traffic Flow and QoS Across the Entire Range of Tactical Dynamics, Network Size, and Network Density
2	The Architecture Will Create the Best Mission Topology Rather than Passively Accepting Network Topology and Routing, as Given
3	The MANET Will Interconnect with Fixed Infrastructure at Multiple, Dynamic Points of Presence Rather than at a Single, Fixed Point
4	The Network Will Create a Distributed Computing Environment where the Applications and Services are Populated/Migrated onto Nodes According to Traffic Flows and Resource Availability
5	The Network Will Have Intelligent Multicast Protocols and Caching Mechanisms to Use Scarce Wireless Bandwidth Efficiently
6	The Network Will Have Cross-Layer Adaptation Mechanisms that Work Together to Optimize Network Performance and Reduce Stress on Inexpensive Physical Layer Devices
7	The Network Will Use Policy to Drive Topology and Load Sharing in the Network
8	The Architecture Will Provide Persistent Caching and Content-Based Access of Information Within the Network
9	The Network Will Support Multiple Network Structures and Multiple Network Frameworks for Delivering High Speed/Low Latency Streaming/Data Services
10	Disruptive Tolerant Networking (DTN) Will be a Native Mode of the Network, not an Overlay
11	New Policies and Policy Controlled Functions can be Introduced Asynchronously, without Code Changes, and Linked Symbolically through an Extensible Semantic Structure

For self-organizing cognitive radio networks to fully achieve maximum benefit, the networking community needs to reframe the question from "*how do we route over a topology?*" to "*What is the right topology to form in order to perform this mission?*" This is a fundamental change in the direction of networking, and is an opportunity that is not present in the conventional terrestrial or more static wireless technologies. For the core network, it is inconceivable that new fiber optic lines could be determined, buried, terminated and provisioned to a network in a sub-second, but that is exactly the opportunity that cognitive radio provides! Similarly, the connectivity of new satellite transponders, or cellular towers cannot be created by a network "on demand", such as a DSA architecture inherently provides. For 30 years, topology has been treated as a given, and the routing as a consequence. With cognitive radio, this relationship can be reversed. The network determines the routing that is needed, and then uses DSA to create the right topology to implement it. The cognitive radio and network community must accelerate research to develop these mechanisms.

A use case for this technology is instructive. Although backhaul bandwidth in much of the world is plentiful, in other regions backhaul is not available, or highly constrained. A cognitive radio network would manage the location of content distributed across end user devices, without discrete servers or any fixed address structure. Many areas of the world would benefit from even localized services that could provide information on locally critical topics, such as medical assistance, food and water availability, and education. Without the requirement for backhaul to fixed services, a content-based cognitive radio network is a feasible and highly desirable outcome of current research. Although wireless may not be the ideal backhaul to distant services, it is an effective means to set up networks whose users are concerned with resources within walking distance more than those on world-wide social networking sites.

An additional consideration is that the cognitive network does not make demands on what is perhaps the scarcest resource in many environments; the intellectual infrastructure to assign frequencies and addresses, designate and configure servers, manage indices, caches and backhaul networks, etc. An information network that operates "out of the box" certainly has appeal and unimagined utility in much of the world, and in many situations[44].

8.5 Infrastructureless Networking

The growth of the networking and wireless industry has been through the continued expansion of the network infrastructure; enhanced technology; convergence towards a common underlying network technology (Internet Protocol) and a focus on new services that can be accommodated and integrated into the network. Certainly, this path has served to deploy an unprecedented amount of capacity and to achieve user acceptance in only a few decades. However, there exist application areas where reliance on infrastructure-based services may not serve as well as some of the pre-existing, displaced technologies. Experience from large scale disasters, such as the Pacific Tsunami and Hurricane Katrina demonstrate dependence on infrastructure creates "common cause" failures. The military is an extreme example of needing independence from infrastructure, but is not unique.

Wireless technology appears to offer a solution to dependence on infrastructure. It does not require fixed landlines; can be operated point to point; does not need hub centers to switch channels, and it is simple to operate, at least in basic form. Any child can use a "walky-talky" or FRS Radio without technical support, help desks, port forwarding

---

[44] For example, in a third world recovery operation where the network was needed to operate immediately in the absence of a large number of Information Technology (IT) specialists, regulatory specialists, frequency assignments, backhaul, and server resources.

policies, static IP assignments and the like. The same is not true for Voice Over Internet Protocol (VoIP)! As the vision of wireless expands from links to networks, the tendency has been to adopt fixed Internet technology and apply it to evolving wireless systems. Self-forming networks have been successfully demonstrated, but an equal focus is needed for service architectures and technologies that exploit the emerging cognitive radio opportunities.

The benefits of minimal infrastructure approaches should not be considered unique to military or institutional uses. Infrastructureless operation also implies that less traffic to and from the core; content is managed closer to users; equipment not only configures itself, it optimizes itself, and services the network offers expand and contract resource footprints as dictated by demand, not static planning. Wireless technology needs to address how networks manage content, resolve names, identify and update routes, and control security. These must be examined from the perspective of the unique constraints and opportunities of wireless communications. Some of the infrastructure services that are candidates for distributed and peer-to-peer implementation are shown in Table 8-2. Many of these services are not currently automated; reducing the human footprint for network operation is as critical as the technical elements.

Table 8-2 Infrastructureless Domains and Approaches

Function	Infrastructureless Approach
Frequency Planning	DSA using common policies in devices to avoid pre-planning frequencies, resolve interference, and manage spectral energy.
Network Planning & Topology	DSA to match spectrum availability and RF propagation with density and topology needs dynamically. Late binding to avoid maintaining IP addressing in dynamic regions.
Index Services	Content metadata advertised locally, and routable metadata descriptions.
Cache Services	All nodes cache content. Content metadata can enable access without full address or URL description.
Information Servers	Any node can store, advertise, and serve metadata defined content.
Name/List Servers	"Late binding" of delivery addresses to address messages and content without knowledge of the destination.
Cell Towers and Land Lines	Maximize reliance on lower cost and expendable nodes to enable density, and create connectivity.

8.6 Networking Implications Conclusion

The advent of Cognitive Radio has generally been approached as being inherently disruptive to the conventional design approaches at the physical layer of the wireless networks. However, the implications of dynamic and adaptive network physical layers are equally, or even more disruptive for networking concepts that are based on operation over static and unchanging topologies and information organization. In the end, it may be that the network layer opportunities of Cognitive Radio are more significant than the underlying physical layer benefits. These benefits can only be achieved if the network concepts are evolved to reflect the dynamics being provided in the physical layer.

Chapter 9 Integrated Model of Cognitive Radio Benefits

9.1 Overview of Cognitive Radio Analysis

The previous Chapters (4 through 7) have outlined specific approaches to the mitigation or exploitation of individual aspects of the cognitive radio environment through individual operational mode decisions. In this chapter, these analyses are integrated into a single model of cognitive radio operation. Some of the performance metrics are specific to prior chapters (such as Probability of Overload) and others are aggregated results from multiple chapters (Mean Noise Floor). The performance metrics that were established as the objectives of this analysis are provided in Table 2-4. The qualitative network performance enhancements provided in Chapter 8 are not included in this analysis.

9.2 Reduction in Hardware Requirements

The opportunity to reduce hardware requirements for devices implementing cognitive radio adaptations can be summarized into four fundamental opportunities:

Reduced Front-end Linearity	Dynamic management of sub-band and filter selection to reduce total energy into the active stages of the front-end.
Reduced Receive Power Consumption for High Reliability Operations	A consequence of the reduction in the front-end linearity is the ability to reduce energy consumption of the entire linear receiver chain, through the final filtering process (typically, the final Intermediate Frequency (IF) stage).
Reduced Transmit Power Requirements for Equivalent E_b/N_0	Reduction in the total power required for equivalent data rates through reduction in in-band interference and adjacent band intermodulation products in the front-end.
Reduction in Spectrum Requirements; Increased Node Density	Increased node density has advantages in reducing the spectrum required to support a given amount of information transfer per unit area, per unit of allocated or shared spectrum. Increased density can lead to lower cost, or more capable services.

These benefits are derivable from the analysis developed in Chapters 3 through 7, as shown in Figure 9-1. Each of these benefits arises from the cognitive radio adaptive mechanisms previously described. In some cases, a single mechanism has multiple benefit areas, for example, the reduction in required IIP3 has benefits in reducing the complexity of the front-end (primarily an affordability benefit) and in reducing energy (an energy savings benefit). Although this discussion is based on the current practice of fixed IIP3 points, the same arguments applies to variable IIP3 architectures, such as have been proposed and used in limited practice [190].

Table 9-1 Component Performance Reduction Rollup from Discrete Function Analysis

Overall Metric		Driven By Functional Characteristic		Eq.
1	Front-end IIP3	Probability of Overload	$P_{\mathrm{Feoverload}}$	(4-15)
		Probability of IMD3 Noise Level	$\mathcal{P}$ (IMD3$_{CR}$)	(4-16)
2	Reduced Receiver Energy Consumption	Probability of Overload	$P_{\mathrm{Feoverload}}$	(4-15)
		Probability of IMD3 Noise Level	$\mathcal{P}$ (IMD3$_{CR}$)	(4-16)
3	Reduced Transmit Energy	Probability of IMD3 Noise Level	$\mathcal{P}$ (IMD3$_{CR}$)	(4-16)
		Noise Floor Difference	S_{TCR}/S_{RCR}	(5-2)
4	Node Density	Capacity (with Spectral Outage)	SIE	(6-14)
		Bits/Area/Hz	C/B	(7-8)

9.2.1 Receiver Front-end Performance Reductions

Chapter 4 described and quantified the opportunity to reduce front-end performance of a cognitive radio and still achieve stated noise floor and front-end overload probabilities, so long as the radio had sufficient spectral coverage to ensure that accessible spectrum was not correlated. This approach has significant savings in cost and energy, and directly relates to the IIP3 reduction relationships shown in Chapter 4. In general, for most practical conditions, the $\mathcal{P}$(IMD3$_{CR}$) (probability of a given density of intermodulation noise) is a driving constraint for establishing front-end design performance. Establishing this value is driven by two considerations; the maximum IMD3 noise that the link can tolerate, and the probability of disruption that would be acceptable to the communications system.

9.2.2 Reduced Receive Energy Consumption

This energy consumed by the analog stages of a dual conversion receiver through the narrowband IF, or analog to digital conversion (typically RF filters, LNA, one or two Local Oscillators (LO) and one or two mixers) is directly a function of the linear range required of the device. McHenry et al. has reported a generalized relationship of the required prime power for a given level of IIP3 design performance [191]. Amplifier efficiency ranges from 30% (typical) to 50% (maximum). A summary of this analysis is shown in Figure 9-1.

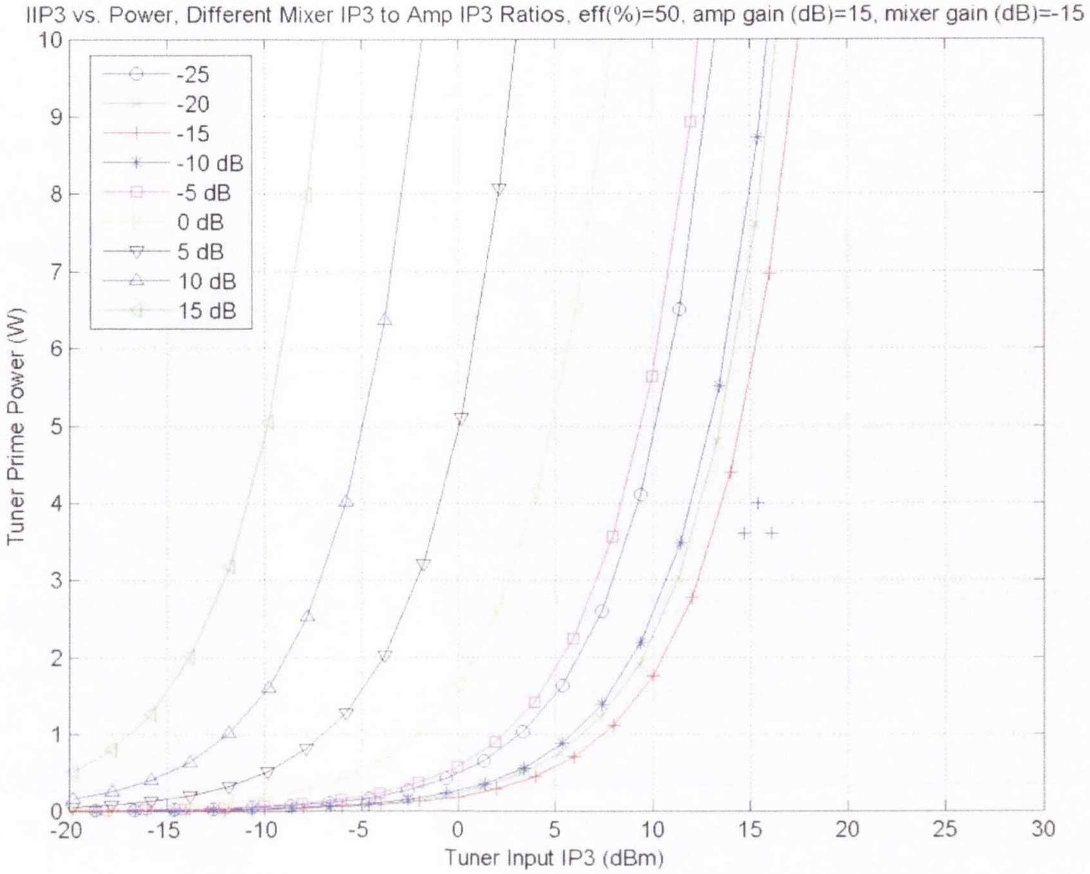

Figure 9-1 ***Tuner Prime Power as a Function of IIP3 and Mixer OIP3 [from McHenry]***

The benefits of reduction in required IIP3[1] values are apparent. Reduction in required IIP3 level has effectively a corresponding linear impact on total energy consumption of the receiver's analog processing chain (prior to digitization, or detection/demodulation).

In many applications, the receiver is the dominant user of energy. Consider as an example a home Wi-Fi hub; operating 24 hours a day at a peak of 54 megabits/second; accessing a 4 megabits/second cable modem for 2 hours a day; and, used at an average of 10% of the modem capacity. These result in idealized payload traffic transmit duty cycle of 0.06%[2]. For such a device, transmit energy is not significant to its total energy consumption, almost regardless of the transmit power! ***While research in communications waveform design has pursued optimizing the performance of transmit waveforms, optimizing these waveforms is a much less important factor in overall energy consumption than managing receiver energy in many of the applications of wireless technology, particularly in short range, variable duty cycle packet networks*** [181].

---

[1] Note that this chart depicts the Output IP3 (OIPT). The OIPT is the IIP3 times the gain (additive in dB).

[2] Compared to a receive duty cycle of 100% for most devices on the receive side of the transceiver.

9.2.3 Transmit Energy Reduction

Transmit power generation by the transceiver has two implications. The first is that the hardware in the device must be capable of developing the required RF energy, and of dissipating thermal load generated through amplifier inefficiency. A second consideration is the device's cost of providing the energy to the amplifier, which is an energy storage and battery lifetime issue; addressed in Section 9.3.3 later in this chapter. The approach described is focused on reduction in the receive noise floor level.

In a non-cognitive radio, the transmit energy requirement is a function of the required link reliability (for a given channel). To examine cognitive radio requirements, equivalent operational reliability must be assured through mechanisms at all layers of the stack, and may enable equivalent reliability with lower amplifier capability.

Transmit energy reduction is obtained through both reduction in adjacent-channel intermodulation induced noise floor elevation, and by better selection of the channel to reduce in-channel noise at the receiver. Assuming that these two processes are independent, the two probability distributions provide the potential reduction in amplifier output power requirements. Figure 9-2 illustrates the reduction that can be provided, depending on the original intended reliability and IIP3 operating point.

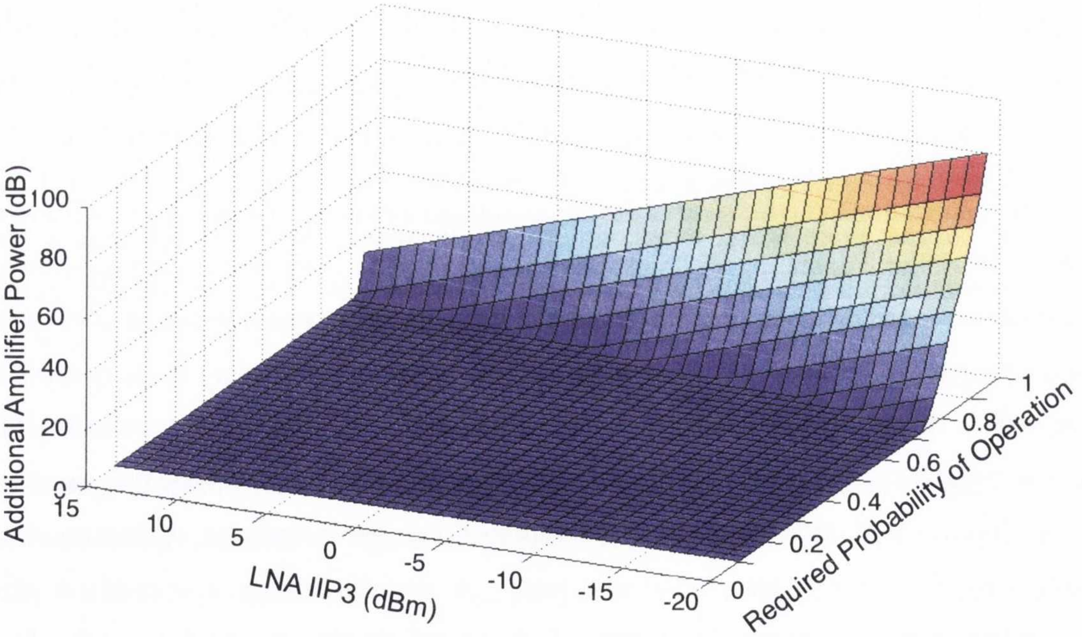

Figure 9-2 ***Reduction in Required Power Amplifier Performance for a Range of Link Reliabilities and baseline IIP3 Levels***

In an extreme case, it would take 80 dB of amplifier gain to make up for the overload induced in-channel noise from a receiver with a -20 dBm IIP3 front-end. At the more typical region of the front-end performance envelope, a design specified to assure high reliability operation would have to provide 30 dB (1000 times) more amplifier output to make up for intermodulation effects, even with an IIP3 of 0 dBm. Note that the benefits are most significant in the case of poor LNA linearity. This example illustrates that significant reduction in required amplifier capability is possible when there is both regulatory and technical flexibility to implement DSA and adaptive techniques in spectrum environments where performance is constrained by receiver non-linear response.

9.2.4 Reduced Spectrum Requirements

Typically, spectrum is not a resource that is provisioned, or paid for incrementally[1] for each unit of use, such as other wireless resources or design drivers: receive and transmit power, linearity or battery storage. However, it is nonetheless an important economic consideration in wireless services; often reflected as a cost allocation to reflect purchased spectrum, an operational constraint due to limitations of allocated spectrum, as additional design complexity or other high performance and cost modes, due to the requirement to maximize use of spectrum resources, such as amplitude modulations required to achieve high bit/Hertz performance[2]. The cognitive radio benefit is the degree to which cognitive radio algorithms can mitigate spectrum shortfalls, and concomitant reduction in the cost (or opportunity cost) of spectrum, in comparison to the complexity that these algorithms introduce, and in comparison to alternative techniques to achieve increased spectral effectiveness.

Chapter 7 demonstrated that high bit per Hertz modulations are ineffective as a strategy to maximize aggregate spectrum utilization in situations where spectrum reuse was a consideration. Figure 6-6 thru Figure 6-8 show that for a typical mobile case, moving from node separation sufficient for low probability of interference in a properly manually de-conflicted spectrum environment to the cognitive radio determined optimal aggregate capacity over a the region can increase node density by approximately two orders of magnitude. Equally important, any outages that might occur are statistically distributed across all users, so no one user has a loss of service, even if the density level exceeds the optimal.

9.3 Increased Cognitive Radio Performance

This section looks at another aspect of the previous analysis; if all other characteristics of

---

[1] Although this regime is contemplated by numerous discussions of spectrum micro-charging.

[2] Which in term require waveforms that are difficult to amplify efficiently, such as those with high Peak to Average Ratio (PAR) in their energy.

the radio are held constant, what is the increase in performance that can be obtained by applying the cognitive radio adaptive techniques discussed previously? There is inherent overlap between these benefits, and the ones cited previously.

Increased Operational Availability	A cognitive radio, with equivalent component performance, should have significantly improved probability of operation, given a sufficient extent of decorrelated spectrum access.
Increased Link Throughput through Lower Link Noise Floor	The noise floor distribution encountered by a Cognitive Radio should be lower than the noise floor distribution of a non-Cognitive radio.
Increased Operating Period/Reduced Energy Storage	The cognitive radio should have a much longer duration of operation with equivalent energy storage, or alternatively, the device can be reduced in mass and volume by reducing energy storage without compromise in operational duration.
Increased Density and Information Capacity	A network of cognitive radios should have increased aggregate density and overall information transfer density per unit area.

Table 2-4 provided the high level metrics by which the improvement in performance of a cognitive radio was assessed. Table 9-2 illustrates how these metrics are derived from the probabilities and performance impacts developed in the earlier chapters, and shows the basis for the aggregation of the individual measures.

Table 9-2 Overall Cognitive Performance Rollup from Discrete Function Analysis

	Overall Metric	**Driven By Functional Characteristic**		
1	Operational Availability	IIP3 Overload Reduction Probability Reduction	$Benefit_{\mathrm{Poverload}}$	(4-18)
		Probability of IMD3 Noise Level	$\mathcal{P}$ (IMD3)	(4-20)
		Single Source Probability of Interference	$\mathcal{P}$ (SSI)	(6-22)
		Pool Frequency Available	A_{spectrum}	(6-10)
2	Link Capacity through Reduced Noise Floor	IIP3 Front End Noise Reduction Probability Reduction	$Benefit_{\mathrm{FENoise}}$	(4-20)
		Noise Floor Difference	$S_{\mathrm{TCR}}/S_{\mathrm{RCR}}$	(5-2)
3	Increased Operating Period/Reduced Energy Storage	Probability of IMD3 Noise Level	$\mathcal{P}$ (IMD3$_{\mathrm{CR}}$)	(4-16)
		Noise Floor Difference	$S_{\mathrm{TCR}}/S_{\mathrm{RCR}}$	(5-2)
4	Node Density & Capacity	Separation of Nodes (Non-CR)	S_{DSA}	(6-8)
		Capacity, with Spectral Outage	*Capacity*	(6-14)
		Spectral Information Effectiveness	*SIE*	(7-6)

Each specific benefit area is discussed in the following sections.

9.3.1 Increase in Operational Availability

The operational availability of a cognitive radio is greatly improved due to reduction in probability of interference through adjacent channel front-end overload and through dynamic in-channel interference avoidance and mitigation. In practice, even if a front-end is not driven to the IIP3 energy level, the link will fail if there is more in-channel intermodulation noise than the link margin allowance provided for, and was not consumed by other causes of non-ideal reception. For most practical conditions, this level is reached well prior to the *IIP3* point, so it is reasonable to assume that the link will fail at least at the rate of $\mathcal{P}(IMD3 > L_{\text{margin}})$ (from Equation (4-16)), when *IMD3*, plus all other link effects exceed the link margin. The link can also fail due to lack of spectrum for a cognitive radio, or improper management of spectrum resulting in an unexpected interference event to a non-cognitive radio. The availability of the non-cognitive radio is significantly impacted by the front-end performance and the degree of acceptable risk in the spectrum deconfliction; while for the cognitive radio it is the much lower probability of front-end overload noise exceeding the margin, and the probability of spectrum pool exhaustion (A_{spectrum}).

The interference probability of a radio is therefore given by:

$$P_{\text{interruption}} = 1\text{-}(1\text{-}\mathcal{P}\ (IMD3 < L_{\text{margin}}))\ (1\text{-}P_{\text{interfere}}) \qquad (9\text{-}1)$$

$$\text{Where } P_{\text{interruption}} \text{ and } \mathcal{P}\ (IMD3 < L_{\text{margin}})) <\!<1$$

$P_{\text{interfere}}$ is the probability of failure of the spectrum management process. This is the probability that a non-cognitive radio finds an interferer, or a cognitive one cannot locate an opportunity. For a non-cognitive radio, $P_{\text{interruption}}$ is the combined effect of the intermodulation probability and the probability of successful spectrum management, as shown below from Equation (4-13).

$$P_{\text{interruptionNCR}} = (1\text{-}P_{\text{interference}})\ I_{x\text{IMD3}}\,(Fe\alpha, Fe\beta) \text{ for } 0 \le x \le 1\text{, where:}$$

$$x_{IMD3} = \frac{\left(\dfrac{BWfc}{k_1 b_0}\right)(PI_{MD3} - 2IIP_3 + k_2) - FE_{\min}}{FE_{\max} - FE_{\min}}\text{, and} \qquad (9\text{-}2)$$

$Fe\alpha$, $Fe\beta$, $FE_{\min}$, and $FE_{\max}$ are characteristic of the spectrum environment

For a cognitive radio, the $P_{\text{interference}}$ term is a function of the size of the pool, and the mean value of the demand, as shown in equation (6-4). Combing this and Equation (4-15) yields:

$$P_{\text{interruptionCR}} =$$

$$\left(\sum_{k=needed}^{Npool} \binom{Npool}{k} (1-duty)^{k}\, duty^{Npool-k} \right) \left(1 - I_{xIIP3}\left(FE\alpha, FE\beta\right)\right)^{PSS}$$

Where : (9-3)

$$x_{IIP3} = \frac{IIP3 - FE_{\min}}{FE_{\max} - FE_{\min}}$$

$$\mathrm{FE}_{\min} < IIP3 < FE_{\max}$$

Both *PSS* and N_{pool} reflect the size of the spectrum pool available to the device. Determination of the binomial parameters was described in Chapter 5.

$$loading = \frac{duty \quad users}{N_{pool}} \tag{9-4}$$

The controlling variables can now be written in terms of the two bandwidth categories (*BW* and b_0) terms from Chapter 3, and the extent of spectrum coverage (f_{low} and f_{high}).

$$needed = users = \frac{loading \quad Npool}{duty},$$

where : (9-5)

$$N_{pool} = \frac{f_{high} - f_{low}}{b_0}$$

Substituting the previously cited references yields the probability of successful operation of both radio designs, as a function of the spectrum made available to them.

From Chapters 4 through 7, the availability probability of a cognitive radio is driven by the flexibility to tune the front-end through a range of frequency selections, and the size of the spectrum pool. The relative availability of the two designs is therefore sensitive to the extent of available spectrum, the relative IIP3 values and the bandwidth of the device. Figure 9-3 illustrates the availability of cognitive and non-cognitive radios as a function of the availability of spectrum. Note that the ratio of mean demand to spectrum availability is constant in this example. If only a few radios are assigned to a spectrum pool, and the number of channels is correspondingly low, the performance of the cognitive radio is well below that of a conventional radio, largely due to the probability that a number of radios will need access to the spectrum simultaneously (5 to 12 channels). With small pool sizes that fit within a pre-selector, the probability of overload for both cognitive and non-cognitive radios is identical, since there is no flexibility for the cognitive radio to avoid interference (5 to 10 channels). As the size of spectrum pool is increased (>10 channels), benefit arises from two sources[1].

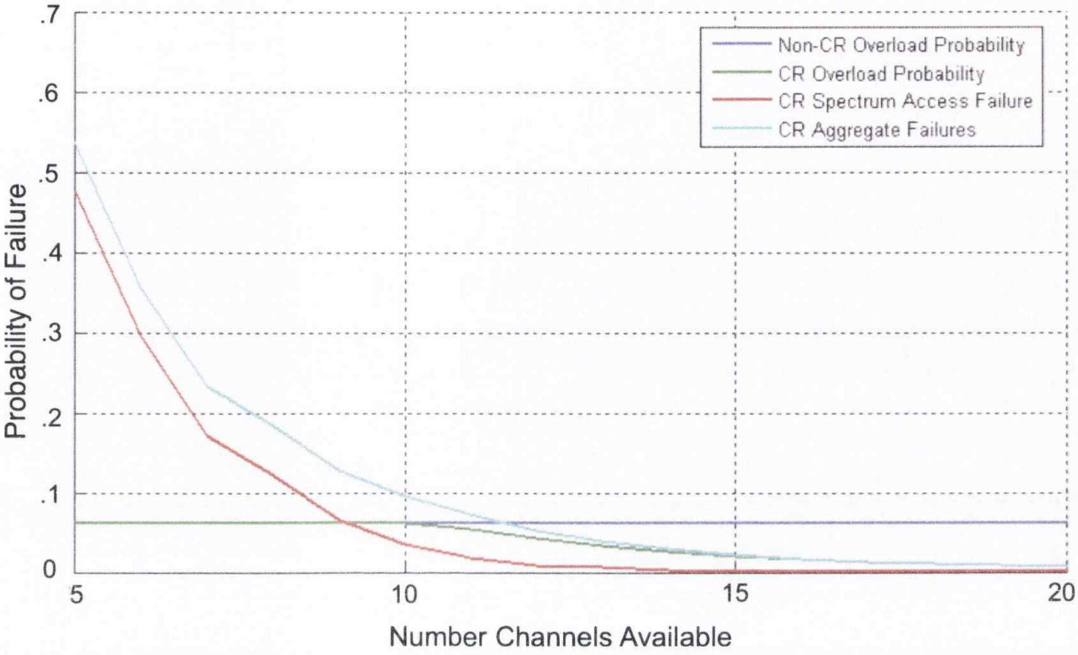

Figure 9-3 Relative Availability for Cognitive and Non-Cognitive Devices for Some Representative Operating Characteristics (Duty Cycle = 10%, IIP3 = -5 dBm)

As discussed in Chapter 4, the additional spectrum choices allows the radio more flexibility in avoiding adjacent channel interference conditions, once more than one pre-selector setting has spectrum available[2]. Increased spectrum also provides benefit in providing the statistical pool shown to be so critical in Chapter 5. This illustrates that increased performance of cognitive radio is dependent on the availability of suitable spectrum options. Without this available spectrum, there is no benefit, and even regret from the implementation of DSA, due to the lack of assured access to spectrum. However, when sufficient spectrum is available, the performance gains are significant, and beyond those that can be readily assured even through additional link margin. The expected value of usage is less important than a statistically significant pool size.

The specific duty cycle is not a fundamental driver in these relationships so long as the pool is statistically large, and above the expected value of the demand. This is a strong argument that all users of cognitive radio would benefit when multiple uses are concatenated, creating options and statistical benefits to all contending users, since one

---

[1] The flat area of intermodulation performance is due to the filter bandwidth being larger than the extent of spectrum, and therefore any signal within this extent has a constant probability of inducing interference, regardless of the number of channels within the filter passband that are available for DSA operation.

[2] In this example, the link was determined to be unavailable whenever the intermodulation induced noise floor exceeded the natural one by more than 20 dB.

large pool is much more desirable than many small ones! This is directly in contrast to the fragmented manner in which spectrum use is currently managed in the exclusive "property rights" model.

9.3.2 Decrease in Noise Floor Probability

Decreased noise floor inherently improves wireless link performance. Section 9.3.1 addressed the availability or reliability implications; in this section, the opportunity to provide increased data rate is analyzed. The notional performance of a relatively heavily loaded transceiver (25% transmit duty cycle) depicted in Table 9-3 is used as a baseline.

Table 9-3 Representative Baseline Transceiver Characteristics

Symbol	Definition	Typical Value
$P_{\text{rcvanalog}}$	Analog Receiver power	200 mw
P_{rcvdig}	Digital Energy Usage in Receiver (A to D and beyond)	100 mw
P_{xmit}	Transmitter section power, not including PA	150 mw
P_{PA}	Power Amplifier output power	1 w
Eff_{PA}	Efficiency of Power Amplifier	30%
xmit	Transmitter Duty Cycle	25%
Range	Pre-selector tuning	1 Octave

Although the actual operation of the channel and waveform may not be at the ideal performance points of the Shannon bound, this analysis assumes that the derivative of the channel bit rate is consistent: i.e., that the relative change in throughput with a change in signal to noise ratio is approximately consistent. This assumption essentially holds link margin and implementation efficiency constant through the range of rates in the neighborhood of the baseline case.

Figure 9-4 illustrates the mean bandwidth improvement possible as function of the cognitive radio adaptation capabilities, using the base case. In this example, the bandwidth increase is weighted by the probability of the intermodulation noise floor associated with the excess energy for bit rate increases. The occupied bandwidth is held constant. It is clear that the poorer the filter performance of the baseline radio, the greater the bit rate improvement that is possible and likely through cognitive adaptation. The probability weighting means that, in most situations, the benefits of cognitive adaptation are not very significant. But, in atypical, but not rare situations, the front-end performance is the driver for the throughput of the device, and cognitive adaptation can make a significant contribution to the aggregate performance and reliability of the link. The benefits of cognitive adaptation are most significant for intrinsically low performance devices. This is another argument that cognitive radio should not be considered solely for sophisticated or high performance applications. Cognitive radio should not be discounted even in lower end products, such as Wi-Fi and consumer networking products.

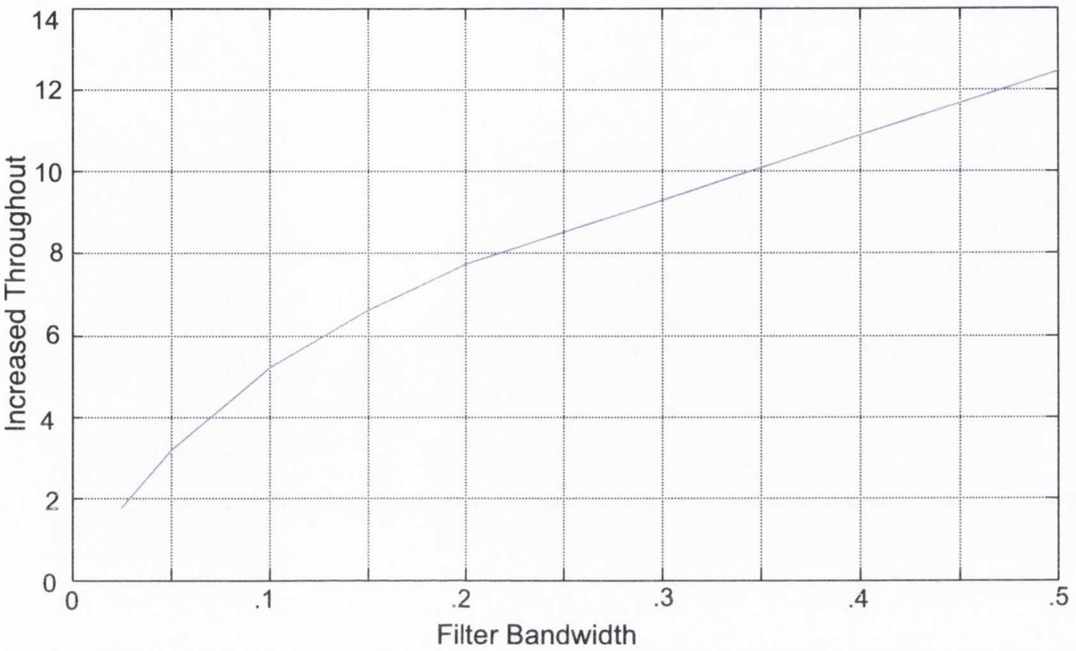

Figure 9-4 ***Bandwidth Improvement as a Function of Cognitive Radio Adaptation Capabilities***

9.3.3 Increased Operating Period/Reduced Energy Storage Mass

The baseline values used above are typical of a small transceiver that might be used as a forwarder or trunk radio in a portable application. A general expression of total energy consumption of this radio is given by Equation (9-6). In this representation, the analog receiver energy, which is strongly proportional to the LNA Output IP3 (which scales with *Gain* times *IIP3*) is a separate term from the digital processing (Analog to Digital conversion through baseband). The digital stages are generally insensitive to the analog stage characteristics, so long as minimum gain to the A to D conversion is achieved.

$$TotalEnergy = P_{rcvranalog} + P_{rcvdig} + xmit\,(P_{xmit} + P_{PA}/\,Eff_{PA}) \qquad (9\text{-}6)$$

Figure 9-5 illustrates the resulting operating life implications for this example case, but is also typical for a range of applications. For high duty cycle rates, the benefit from receiver energy reduction is not as significant (since transmit energy usage dominates), but as the duty cycle decreases, the benefits of receiver energy reduction increase, and becomes very significant in the region below 35%. As a point of comparison, the Wi-Fi example cited previously has a 0.0067% utilization of the Wi-Fi device over a 24 hour period! In practice, current wireless LAN MAC layers rarely achieve over 50% transmit duty cycle even when fully loaded.

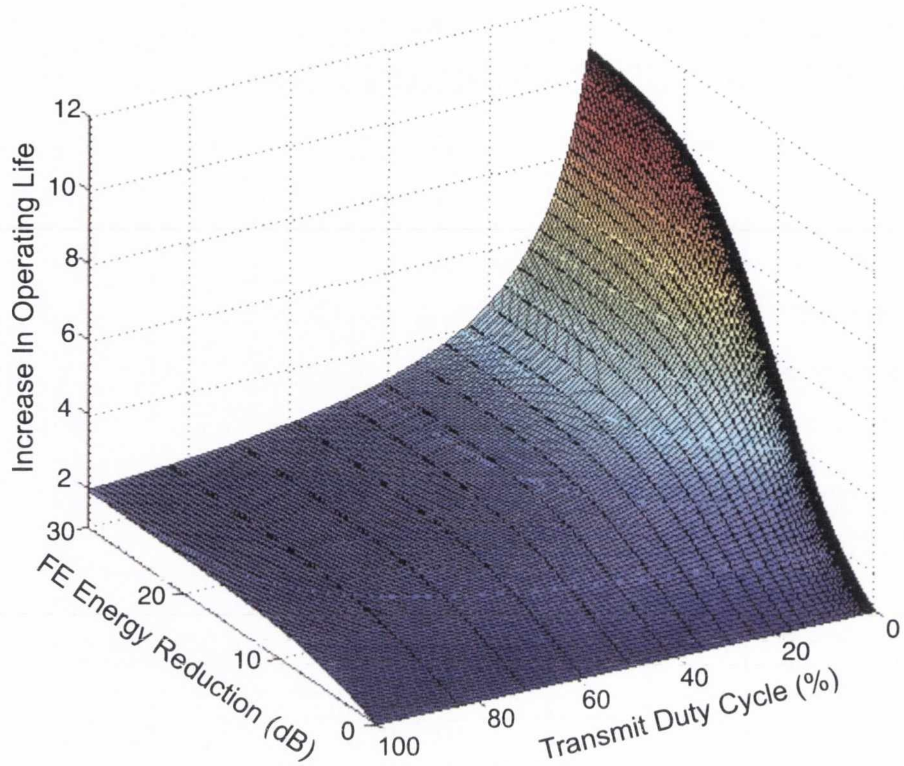

Figure 9-5 ***Operating Life Extension from Reduced Front-end Energy Consumption for a Range of Duty Cycles***

9.3.4 Increase in Aggregate Capacity

Section 9.3.2 addressed the link capability improvements possible through the inclusion of cognitive radio technology. Aggregate capability is the sum of the throughput of all of the nodes within a single region, as analyzed in Chapter 5's Spectral Effectiveness discussion. From Chapter 5, it is clear that the transition from a non-interfering to an interference tolerant environment creates orders of magnitude increases in aggregate capacity, as shown in Figure 6-8 and Figure 6-10.

9.4 Fungibility of Benefits

The previous sections described the benefits of cognitive radio from the perspective of using the performance benefits of adaptation to specifically reduce resource requirements, or purely to enhance performance by increasing availability and capacity, and reducing energy and noise. Alternatively, the benefits can be distributed among multiple objectives, and allocated across these four dimensions. Although previous sections have discussed these objectives as if they were orthogonal axes, they are surfaces on which a radio design can operate.

As an example, the possible reduction in required IIP3 performance can be traded for increased reliability or probability of intermodulation induced noise floor, as described in

Sections 9.2.1 and 9.3.2. One possible allocation of benefits could simultaneously achieve a 99% reduction in front-end energy (20 dB), and reduce intermodulation noise (90% point) by 20 dB simultaneously for filters of 20% bandwidth, as shown in Figure 9-6. These relationships are provided by Equations (4-17) through (4-20). Note that cognitive radio benefits are not monotonic, front-ends with either high or low performance filters have greatly reduced benefits, and low performance filters provide few options for cognitive adaptation, and thus, greatly reduced aggregate benefit.

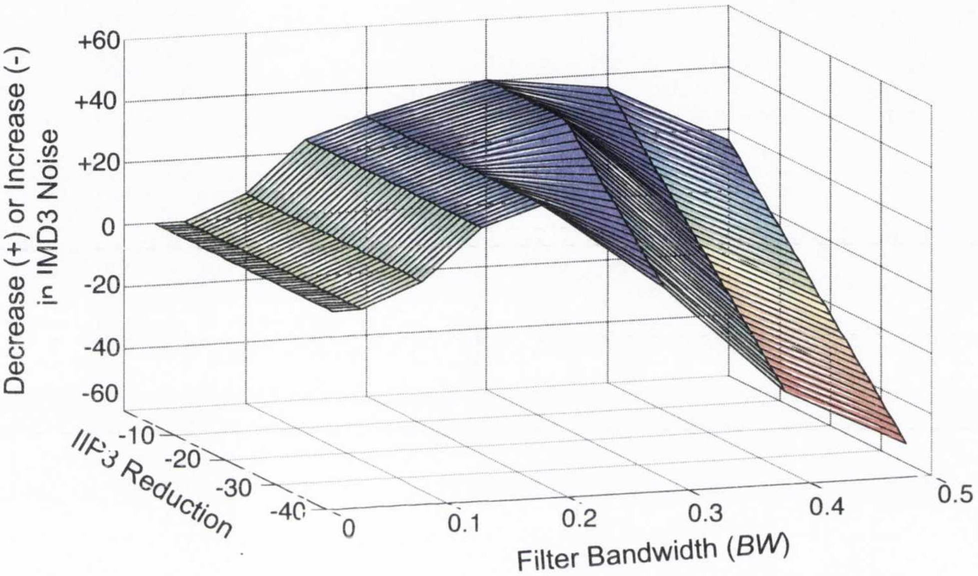

Figure 9-6 ***Trades between Front-end Linearity Reduction, IMD3 Induced Noise and Bandwidth***

This figure indicates that with very high performance filters (typically, pass-band under 5%) there is little mean energy savings due to reduction in the noise floor with these filters. The benefits can best be exploited profitably by the reduction in the required IIP3 level. As the filter bandwidth increases (decreasing performance), the choice between performance and component reduction becomes more complex. For example, with a 25% filter, up to 25 dB reduction in IMD3 noise (90% case) can be achieved, or the IIP3 can be reduced by 30 dB, with neutral performance impact. Between these two extremes, a 20 dB reduction in IIP3, results in only a 6 dB reduction in benefits (to approximately 19 dB IMD3 noise reduction at the 90% occurrence probability). Selection of the optimal point within the trade space is system and application specific.

9.5 Additional Resources Required for Cognitive Radio Adaptation

The benefits described in the previous section require, as a minimum, the introduction of cognitive adaptation algorithms within cognitive radio nodes; and in most cases, some hardware specific to the environmental sensing function. Potentially, these additional functions that may compete for node or platform resources. Potential platform resources fall into four areas:

Environmental and Spectrum Sensing	The additional functionality required to perform spectrum sensing of total energy in operating bands (pre-selector options), and to determine the availability and potential interference of specific signaling channels.
Device Performance	Specific performance thresholds that cognitive adaptation may impose on elements of the wireless device that would already have been required for non-adaptive operations.
Control and Algorithm Processing	The additional processing resources (processing and memory) that must be provided to implement the algorithms introduced by the adaptive techniques described.
Additional Communications for Awareness	The additional communications required to support any collaborative decision making that is not local to the individual nodes; such as selection and communications of the signaling frequency, bandwidth, waveform options, etc.

Each of these considerations is described in the next paragraphs.

9.5.1 Environmental and Spectrum Sensing Resources

A core element of the cognitive radio is the ability to sense and adapt to the spectrum. Typically, this is assumed to be digitization of the pass-band, followed by Fast Fourier Transform (FFT) to the frequency domain. Another possible implementation that could be applied to a more limited performance cognitive radio is to use one or more of the receiver channels as individual channel tuned sensors that can be scanned through individual frequencies and the total energy in each one sensed through the AGC or Received Signal Strength Indication (RSSI) capability that is generally available in most radio platforms. In this case, the sensing interval is driven by the tuning time (typically driven by the settling time of the local oscillators) and the time constant of the processes measuring the signal (AGC or RSSI). Some reported 802.22 approaches provide unconventional, unique, but waveform specific analog implementation of TV signal and "whitespace" detection [41]-[42]. Cabric and Brodersen describe some of the unique challenges in implementing the sensor function of a cognitive radio with limited dynamic range [192].

The DARPA WNaN represents one of the first purpose-built cognitive-radio designs. The control process in the node periodically tasks one of four symmetric transceivers to perform the sensing function, interleaved with communications tasks. This eliminates the requirement for dedicated sensing hardware, and still provides high-speed, high-resolution sensing that can be interleaved with other duties. Such a strategy is very consistent with both TDMA and CSMA channel access schemes. Additionally, this approach enables the cognitive radio to balance and dynamically trade the benefits of more sensing capability with more communications capability based on the environments in which the radio operates. In TDMA or MIMO architectures, it is also possible to apply multiple transceivers to sensing simultaneously, greatly increasing the utility of short sensing intervals.

9.5.2 Digital Processing and Storage Requirements

The substitution of digital logic for analog complexity and performance is central to the concept of cognitive radio. The presumption of that argument is that the implementation of the digital processing is less costly than the utilization of approaches that stress the constraining resource (s) (power, energy, thermal dissipation, linear range, spectrum, etc.). It is assumed that the logic described in the earlier section will be implemented in a digital processing model. Non-digital models have been proposed [193], but none are near the maturity that would enable them to be considered as alternatives for implementation of a cognitive radio. Therefore, the standard models of digital processing are applicable, such as General Purpose Processors (GPP), Digital Signal Processors (DSP), Field Programmable Gate Arrays (FPGA) and Application Specific Integrated Circuits (ASIC).

There are several processing models for implementation of cognitive radios. The conventional expectation is for programmed instructions describing the operation of the radio in a procedural description, such as in JAVA, C, or C++. Recently, there has been advocacy for non-procedural models of cognitive radio policy expression, such as Ontology-based and rule languages, such as Ontological Web Language (OWL) and OWL+RULES [194], [195]. Feeney and Lewis show that policy languages can be extended to recognize, and address policy development by distributed authorities [196], [197].

Unfortunately, these technologies are too immature to have an accepted baseline for estimating either the total number of rules required, or the execution resources associated with them. Successful experimental implementations of OWL-based reasoners in cognitive radios have been reported, but have not been experimented or reproduced by the community [198]. The Lehigh University benchmark has made progress in establishing metrics for reasoning speed and throughput, but has not been extended for rule and constraint solving [199].

Table 2-3 and Figure 2-4 provide a high level analysis of the cognitive radio process chain and are summarized in Table 9-4.

Table 9-4 Processing Functions and Models

Function	Processing Model
Sense Broad Energy by Band	Aggregate energy detection can be performed using low resolution FFT techniques, or depending on the filter, just measure RMS energy. Resolution needed is limited to the filter bin size, and no resolution is needed within the filter pass-band.
Determine Impact on Radio	Computationally trivial, based on estimate of intermodulation noise floor using estimating techniques in Chapter 4.
Select Initial Candidate Bands for Analysis	Computationally trivial, based on ranking of predicted intermodulation noise floor.
Sense Frequency Occupancy	This function requires more extensive frequency resolution, down to the lower of the resolution of the protected signals in the band, or the bandwidth of the cognitive radio. In lower VHF, this implies a resolution bandwidth of from 6.25 to 25 kHz. In bands that are dominated by more wideband signals, the resolution requirements can be relaxed appropriately. FTT processing is typically provided by non-GP processors, such as FPGA, ASIC or in lower performance applications, DSPs or GPPs. FFT algorithms and associated processing resources are well documented in the literature.
Assess Candidate Frequencies for Opportunities	Comparing measured power with a spectrum reuse policy, and then providing relative ranking of candidate frequencies is not intrinsically consumptive of computational resources. It should be noted that there are significant benefits to a language (rather than code-based) approach to this functionality. In this case, processing resources are a significant constraint.
Assess Footprint Impact on Other Users	The determination of the footprint is computationally trivial. The significant assumption of this processing is awareness of the range of propagation exponents that are encountered This awareness can be created by comparing received signal strength transmitted energy, and distance between the nodes.
Noise Floor Analysis	The information needed for noise floor analysis is provided inherently by the spectrum sensing function. The number of candidate frequencies is limited, and the decision criteria are simple.

It is clear by the simplest inspection that the functions (with the exception of MIMO processing) are computationally trivial in comparison to tasks such as waveform, FFT and coding/decoding processing, so the resources required to perform them are not significant constraints in considering these implementations. Again, each of these steps is depicted as if they were sequential, but this assumption is not necessary, as each of these processes is, in general, composable and severable.

9.5.3 Additional Communications for Awareness

A cognitive radio may not be able to sense all the information required to make the decisions necessary to operate optimally in all environments. One area where this is certainly true is in spectrum policy, where some uses of the spectrum are not active, and therefore not apparent to sensing-based algorithms. This awareness must be obtained and retained in the radio. This process is not without resource impact, both for the collection and retention. Some of this information can be prepared and stored in the radio in advance, while other information is situational and dynamic, and so must be collected periodically at a rate that exceeds the rate at which the underlying process changes. It can be assumed that information items in the first category (static) can be downloaded to the radio at deployment, or periodically as part of a software update process. Since these operations can be scheduled in such a manner that they do not use critical resources, they can be considered to have minimal marginal cost. The more complex case is composed of data that constantly changes (dynamic) and must be collected and communicated as dictated by the process involved, regardless of the cost to the wireless device or network.

There has been very little published work exploring the likely size of regulator provided policy databases, or how compact these databases can be when stored in the executable format. Both Shared Spectrum [75] and SRI [74] have reported work on ontology-based reasoners for spectrum policy. McHenry provided an estimate of the total size of the policy database used in their experiments [200], and also estimated that an operational prototype would require three times the policy statements. The complete database for all of these is small compared with the storage of even the simplest processing device. Much of the information required by the decision process can be implicitly detected by receivers without direct and control channel communications, such as pointed out by Sutton et al. [185] using cyclo-stationary information as a self-signaling mechanism about bandwidth choices.

9.6 Scaling to Larger Deployments

Another issue that has received little research attention is the scaling effects of cognitive radio control algorithms. However, if scalability of cognitive radio cannot be assured, then its viability as a basis for future networking is suspect. Analysis by Doer et al. [201] shows that the scalability of a central spectrum server is limited by a growth in inter-node traffic that scales with n^4. Clearly, such approaches are not scalable to Internet scale, or even intranet levels. Figure 9-7 illustrates the conventional view of the deployment of

Mobile Ad-hoc Networks (MANET) and might incorporate DSA. In this model, the network rendezvous all of the nodes onto a single frequency.

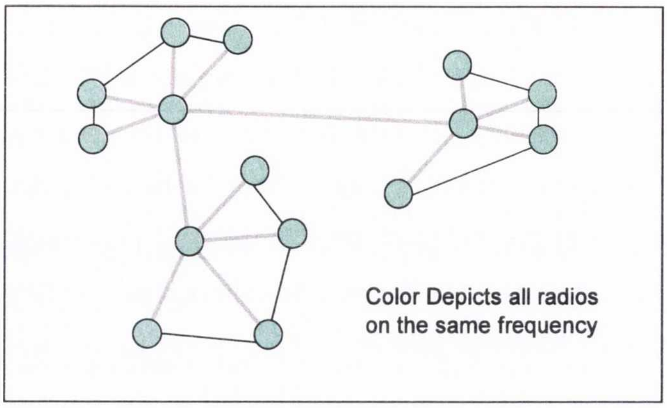

Figure 9-7 Typical MANET Network

This architecture has the severe scaling limitations pointed out by Doerr [201]. In WNaN, the author proposed that providing multiple separate transceivers, each forming small local networks would be one mechanism by which the network could avoid these scaling constraints. Small, local networks on separate frequencies (made possible by DSA) could decouple the effect of changes in one portion of the network to reduce the thrashing.

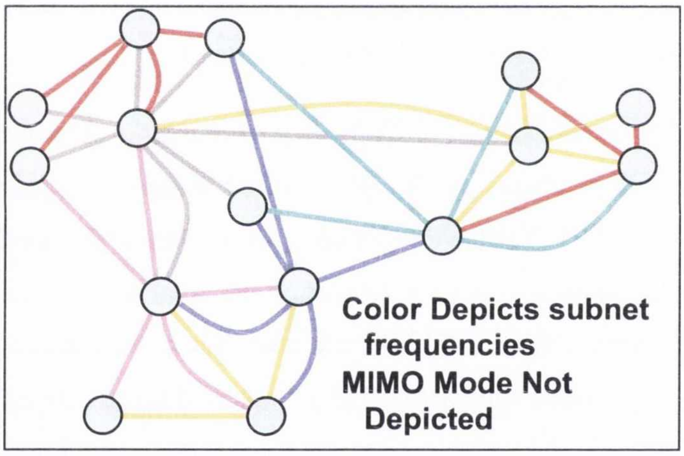

Figure 9-8 Cognitive Radio Enabled WNaN Network

Figure 9-8 illustrates the network, which exploits the spectrum adaptation to both provision spectrum, and to enable networking scaling (in extent and density). In this architecture, each node connects to a different set of neighboring nodes through multiple frequencies and autonomously managed sub-networks. Each frequency sub-network is independent of its neighbors, and relies on routing rather than radio link range for longer distance communications.

This architecture is enabled by DSA and cognitive radio technologies. The manual and static assignment and management of this many frequencies in a complex and constantly changing network would be impractical, if possible at all. These approaches are only practical with automated management and interference control. The second enabling consideration is the use of the adaptation of a cognitive radio to use lower cost and performance transceiver components in each of the nodes, making the multi-transceiver configuration affordable at a lower cost than if a single, high quality transceiver was utilized. Not only does the cognitive radio functionality improve the link-level performance; it is the key factor in affordably replicating the physical layer, by reducing performance requirements, and thus cost of each of the replicated transceivers. This enables the wireless architecture to leverage one of the original, and fundamental principles of the Internet; creation of reliable communications through diverse routes, rather than highly reliable links.

This structure defers decisions regarding optimal size of DSA connected sub-networks to the actual operation of the network. It has the option to partition single frequency DSA sub-networks to reduce the overhead in managing them, but at the cost of introducing more routing hops. Alternatively, it can coordinate larger DSA sub-networks, at the cost of additional overhead, but reducing the hops to transverse the network.

This structure provides more than multi-transmission capability [202]. Each of these DSA sub-networks is a unique management domain that can operate on identical, similar, or even distinct policies. For example, bandwidth on a 900 MHz DSA network might be quite rare, and thus used as a last resort, while bandwidth on a 6 GHz DSA network would be plentiful, but require more hops. Just considering the path losses, this insight to balance the spectrum and routing is an integral problem. Network topology approaches, such as Hazy-Sighted Routing, reported by Redi, can abstract the details of more distant sub-networks to reduce network routing information exchange [203].

9.7 Non-Recurring Software Development Cost Considerations

The discussion in this Chapter has generally focused on the benefits from reducing the performance requirements, and thus, the recurring cost, of the components within the cognitive radio. This section considers the costs that do not recur on a unit-by-unit basis, but which are expenses in establishing and maintaining the software to implement the behavioral algorithms described previously.

As an example, if the additional software for a cognitive radio (such as described in Chapters 4 through 8) is on the order of €10M, then this investment would only have to reduce component cost by several euro on a fraction of the large volume of wireless devices that are sold world-wide, to both recoup the investment, and provide return on the investment. Ideally, the software would be sufficiently general so as not to be made obsolete by evolutionary component technology advances. Since over 314 million consumer home wireless devices are predicted to be sold in 2010 [204], even a small portion benefiting from this technology would be sufficient to make the economic case for savings far in excess of that needed for cost recovery! In fact, the algorithms in these chapters are all expressed in relatively simple, but sufficiently formal descriptions to argue that the implementation of the required functionality is relatively straightforward. Chapter 4, for example, describes selection of operating frequency bands in structures consistent with Predicate Calculus notation. This would argue for reasonable implementation efforts, since these logical structures map well to classical procedural code. In most cases, the algorithms described in all but Chapter 9, are generally stateless, or should perform sufficiently well in a stateless implementation.

In applying these techniques to higher levels of the communications and networking stack, the amount of interaction, and the inherent quantity of process states involved, does increase. In fact, the complexity of a fully "stack smashed" Cognitive Network Device is quite daunting. A good argument could be made that the implementation of these logical control systems should not be pursued as an increasingly complex software development effort; but instead, would benefit from implementation techniques provided by "Artificial Intelligence" derived tools, such as; rule-based constructs as in PROLOG [205], its descendents, such as Transaction Logic [206], Frame Logic [207], and the Semantic Web derivative SWRL [208]. All appear well suited to be the basis for cognitive radio implementation. Previously cited work describes some of these implementations in cognitive radios and applications. In particular, this work demonstrates that a spectrum logic can be created without the use of a negation operator, which otherwise complicates the logic of the reasoner component[1]. Progress on defining a common cognitive radio language has been reported by industry and academic participants [209].

---

[1] This structure allows rules to assert that an emission is permitted, and also to assert that an emission is precluded. A meta-rule then combines these (almost) complementary concepts by only allowing an emission if it is both permitted by at least one rule, and not precluded by any rule. Therefore, it is possible, and not a violation of logic, for an emission to simultaneously be permitted by one rule, and yet be precluded by another. As a concrete example, a band of frequencies may have emissions permitted by a general rule, but another rule could specifically constrain a specific frequency and cause it to be precluded.

Chapter 10 Conclusions

10.1 Overview

The preceding chapters have established that the introduction of cognitive radio technology provides a number of quantitative benefits and enhanced functionality that cannot be provided by conventional wireless designs. Cognitive radio technology provides the capability for quantitative improvements in the performance of wireless devices, mitigation of some inherent component limitations that otherwise constrain the performance and operating characteristics of wireless devices and networks, and the flexibility for adaptation in the higher layers of the network.

The advantages of cognitive radio and DSA do not just derive from increased mean values of performance: they provide significant improvements in the reliability of wireless systems. While technologies are often measured by their mean values of performance; it is the reliability of that performance that is often more important. For example, how many customers would select a cable modem speed that was increased by two, although with only 90% availability? The mean value would be increased, but for most users, its desirability would be decreased. The contribution of cognitive radio should be examined not just from the perspective of mean performance, but also from the perspective of performance at various (stressing) levels of reliability.

Six overall conclusions arise from the work reported in this dissertation.

1. The environments of a cognitive radio can be described by a set of closed-form probability distributions that are directly derivable from spectrum measurements, and closely approximate the measured environmental characteristics. These distributions can be prepared for both strong and weak signal characterizations, and can also be synthesized to interpolate between, or extrapolate from, the measured environments.

2. These closed-form environmental expressions lead directly to estimates of the environmental sensitivity of the performance of cognitive and non-cognitive radios by providing probability distributions of environmentally driven performance metrics. A discrete analysis of spectrum samples can only show the performance of techniques in individual environments; the use of the closed-form relationships enables cognitive radio functionality to be "proven" over a range of environments.

3. A relatively straightforward set of environmentally informed decisions can significantly improve the performance of wireless devices in a DSA regime. This performance benefit can be realized through a mix of benefits: higher mean performance; greatly increased probability of successful operation; reduction of hardware performance levels, enabling lower cost and energy consumption wireless devices, and through a mix of each of these benefits. While DSA was initially proposed to address spectrum access shortfalls, it can be a fundamental tool to address other, equally or more important, aspects of wireless and network performance.

4. The use of DSA changes more than spectrum management: by enabling devices to adapt to their environment, it changes the optimal approach to maximize aggregate spectrum utility. Active management of network spectrum re-use requires fundamentally different strategies than those appropriate to minimize link spectrum usage. These strategies are dependent not just on spectrum awareness, but awareness of the propagation environment of the network and t0 the potential recipients of interference.

5. While non-interfering spectrum sharing with non-cooperative devices is a worthy objective for DSA, significant additional improvements in performance and density are possible when the transition is made to spectrum sharing regimes in which all users of the spectrum can assume that the other devices are capable of interference tolerant operation through their own adaptive behaviors. By mitigating the consequences of interference, each participant in a spectrum sharing regime can pursue higher interference risk strategies, in order to maximize spectrum utility. A high degree of interference may be present when the spectrum is being optimally utilized, as in some MAC layer protocols. Cognitive radio spectrum sharing with other cognitive radios is significantly more beneficial than cognitive radio sharing with conventional radios.

6. Cognitive radio, with DSA functionality, is also enabling of significant adaptation in the upper layers of the network; which can itself be a source of future performance benefits. The ability to adapt the physical layer in response to the electromagnetic environment is only one aspect of the cognitive radio benefit; the other aspect is physical layer adaptation in response to the upper layer environment, including traffic flow, routing needs, and content flow. This level of

adaption can enable topologies that are responsive to the dynamics of the traffic flow, rather than link closure.

10.2 What if We Succeed: Future Research Needs

Much of cognitive radio research is focused on current problems in wireless technology, and clearly it is essential that cognitive radio offer useful solutions to these problems. However, it should also be recognized that when cognitive radio deployment occurs, a consequence will be new and emergent issues. The deployment of tens of cognitive radios will have little impact: the deployment of hundreds of thousands will change the nature of the spectrum environment, create new issues in network operation, and require policy control solutions that are beyond those envisioned today. In particular, the inclusion of dynamic and opportunistic spectrum in cognitive radios will fundamentally change the nature of a radio deployment, due to significant reduction in the constraints of spectrum availability and management. It is important that research looks ahead to this epoch, since there is little benefit in removing one technologic obstacle only to immediately encounter another. The following discussion is intended only to introduce needed research topics.

(1) How Do We Make Provable Assertions About Cognitive Radio Performance?

A particular need that cognitive radio will accelerate is the transition from reliance on wireless networking practice to reliance on wireless network theory. While current wireless networks are either conceptually simple enough (such as access points, cellular and point-to-point links) or sparse enough (military devices) to make discrete techniques (such as simulation) practical, future deployments offer the potential to create density and interaction challenges beyond those currently encountered in wireless practice, and more complex that the hierarchical and decoupled autonomous systems model that is the core concept of the Internet [210]. A more theoretical foundation for wireless network research appears to be fundamental to dealing with the challenges that will arise.

(2) How Can Optimization Be Performed Across All the Layers, and Recognize all of the Environmental Interactions Among Cognitive Radios?

Routing alone is a hard problem to solve, even with stable traffic models and topology, and non-real-time requirements. With cognitive radio, the decision is real-time; the traffic is highly unstable; interaction among the nodes occurs at multiple layers (from spectrum interference at the physical layer, to transport layer error correction); and the topology itself is constantly changing. Even if the requirement for optimality is relaxed to a criterion that only approaches optimality, the complexity is beyond what is currently addressed with the wired Internet. Heuristic approaches may point to low processing demand solutions; but the cognitive radio community is in a position where it is about to create more alternatives for a device than can be processed by the algorithm methods currently available.

(3) How Can Cognitive Radio Algorithms and Reasoning Be Expressed?

It is highly likely that some form of cognitive radio will see operational use within a decade; if not much sooner. If left to natural evolution of the field, the description of behavior will be provided through the current computer science techniques, which largely rely on procedural methods. The likely result of this will be islands of technology capabilities, each targeting specific problems, but not readily integrated with other solutions to other problems.

SCC41 has established a policy language as one of its goals. The SDR Forum has also reported progress on development of a joint industry-academic community project [194], [209]. Without a higher level description of the "reasoning" of a cognitive radio, the community will find itself where much of the current software community finds itself: with complex and unintelligible systems that are hard to maintain, insecure, and lack visibility and insight into their fundamental principles of operation. Cognitive radio would appear an application that would benefit from alternative computer sciences, such as declarative and rule-based structures.

(4) How Can Cognitive Radio Tools Integrate Learning and Reasoning?

The reasoning engine is only one half of the tool set that must be integrated into a cognitive radio. There is an equal need for learning-based engines that offer so much promise for the endogenous portion of the decision process. Integrating these two very different structures into a cohesive theory and implementation that can be placed into devices is a significant challenge. Some concepts need to be inherent (declarative) in the systems that will be constructed, but others are best learned through experience. This hybrid is a technology that does not exist in mature enough form to be the basis for wireless exploitation; but is needed.

(5) How Do We Assure Cognitive Radio Stability?

The logistics of acquiring, programming, and operating wireless nodes has generally limited the size and extent of the networks that have been used to prototype cognitive and ad-hoc wireless technology. When the node count is in the hundreds, each node can be treated as an individual, and the dynamics of the system is sufficiently limited that it can be analyzed by discrete tools. When the wireless element was just a one-hop access path (Wi-Fi, Cellular), then the Internet backhaul decouples the interactions amongst them. But, in contemplating networks that have large membership, not in the hundreds, but in the hundreds of thousands of interacting nodes, explicit understanding of the behavior of individual nodes probably becomes unachievable, and a transition to more analytic and systemic treatment of network behavior is mandatory. The first question referred to the need for analytic methods to determine performance; an equally important need is for analytic means to assure stability.

Density will create issues at both the network and the physical layers of the device. To some extent, the dense cognitive radio network will function like an ecosystem, similar to a population model where it is difficult to make assertions about the future state of individual organisms, but where the aggregate dynamics and end state of the population as a whole can be predicted with high accuracy.

The need for proving stability of extremely large number of nodes would argue for a transition from discrete simulation to formal analytic models of the dynamics of systems of nodes. This transition might have significant implications for how the design of such large scale and highly coupled systems is approached. Formal proof of stability may have to become a fundamental driver in MAC, network and application layer design. Appropriate damping of impulsive modes, excitation resistance, randomizing responses, and varying resonant modes, may need to become integral in the design of these layers. The types of design considerations normally associated with a structure, or a control process may have to be of equal or greater importance than today's concerns with throughput and latency.

(6) How Should Decision Theory Be Integrated into Cognitive Radio Algorithms?

Classically, radios obtained the information they need as an intrinsic byproduct of their operation. There was little need to define information-seeking behaviors. However, a cognitive radio will have a large number of choices to make: to obtain, defer, or ignore potential information regarding its environment. There has been a large and robust literature regarding the choices that a cognitive radio should make in selecting its operating modes, but little attention has been paid to another dimension of its decisions: how does it balance the resources needed to obtain information, with the benefits it derives from having that information. This is a problem in Decision Theory; balancing the cost of obtaining information against the benefits arising from better decisions that are based on that information. A truly comprehensive theory of cognitive radio must address this, along with the more conventional electrical and network engineering decisions.

10.3 Overall Conclusions

Cognitive radio research has now advanced to the point where a number of techniques and technologies have shown desirable experimental results. Some of the technology has advanced to the point of reporting meaningful field test results that strongly argue that cognitive radio can begin to fulfill its promise to revolutionize wireless communications.

A significant advance in the attractiveness of cognitive radio technology can occur if the functionality of sensing and decision-making implicit in the concept could be considered to be substitutional, rather than supplemental to the constraining (cost and performance) elements of a conventional radio. In this manner, equipment and service developers can make direct engineering and cost trades between achieving performance and reliability through robust component performance and spectral resources, and through adoption of cognitive radio methods. Such an approach avoids the necessity for an absolute cognitive/non-cognitive partition, and creates a unified engineering approach to synthesize the benefits of both conventional and emerging approaches to high assurance and affordable communications.

There would appear to be sufficient technology and experimental platform development underway to enable early deployments of cognitive radio platforms, probably leveraging DSA as the initial rationale for their deployment. The community should begin to transition its vision from cognitive radio as an adjunct to the current family of wireless devices and services, to a focus on the new environments, opportunities, and capabilities that can be created uniquely through the deployment of cognitive radio.

Appendix A Acronyms and Other Terms

ACK	Acknowledgment
ADC	Analog to Digital Converter
AGC	Automatic Gain Control
AI	Artificial Intelligence
AODV	Ad hoc On Demand Distance Vector Routing
ARM	Active Reliable Multicast
ASIC	Application Specific Integrated Circuit
AWGN	Additive White Gaussian Noise
BAA	Broad Area Announcement
BER	Bit Error Rate
BGP	Boundary Gateway Protocol
BLAST	Bell Labs Adaptive Space Time Coding
BPSK	Binary Phase Shift Keying
CDF	Cumulative Density Function
CDMA	Code Division Multiple Access
CER	Cost Estimating Relationship
CFP	Connection Formation Probability
CMNI	Consider Marginal Noise Impacts
CMOS	Complementary Metallic-Oxide Silicon
CR	Cognitive Radio
CSMA	Carrier (Collision) Sense Multiple Access
CTS	Clear to Send
DARPA	Defense Advanced Research Projects Agency (US)
DFS	Dynamic Frequency Selection
DFT	Discrete Fourier Transform
DOS	Denial of Service
DSA	Dynamic Spectrum Access
DSP	Digital Signal Processor
DSSS	Direct Sequence Spread Spectrum
DTN	Delay Tolerant Networking
DTN	Disruption Tolerant Networking
DYSPAN	New Frontiers in Dynamic Spectrum Access Networks Conference
E_b/N_o	Energy per Bit over Noise
ECC	Error Correcting Code
EDAC	Error Detection and Correction
EHP	Equivalent Hardware Performance
EIRP	Equivalent Isotropic Radiated Power
ELB	Emergency Locator Beacon
EVM	Error Vector Magnitude
FDMA	Frequency Division Multiple Access
FE	Front-end

FFT	Fast Fourier Transform
FPGA	Field Programmable Gate Array
FRS	Family Radio Service
FSK	Frequency Shift Keying
GPP	General Purpose Processor
HF	High Frequency
HTML	Hyper-Text Markup Language
IETF	Internet Engineering Task Force
IF	Intermediate Frequency
IFFT	Inverse Fast Fourier Transform
IIP2	Input Intercept Point, Second Order
IIP3	Input Intercept Point, Third Order
IMD	Inter-Modulation Distortion
IP	Internet Protocol
IP	Intellectual Property
IPSEC	Internet Protocol Security
IRG	Internet Research Group
ISI	Inter Symbol Interference
ISM	Industrial, Scientific, and Medical Band
IT	Information Technology
ITU	International Telecommunications Union
LAN	Local Area Network
LB	Locator Beacon
LBT	Listen Before Talk
LNA	Low Noise Amplifier
LO	Local Oscillator
MAC	Media Access Control
MAN	Metropolitan Area Network
MANET	Mobile Ad-Hoc Network
MEMS	Microelectromechanical Systems
MIMO	Multiple Input/Multiple Output
NET	Network Layer
NIB	Non-Interfering Basis (or Not to Interfere Basis)
NRAO	National Radio Astronomy Observatory
NSF	National Science Foundation (US)
OFCOM	Office of Communications (UK)
OFDM	Orthogonal Frequency Division Multiplexing
OIP3	Output Intercept Point, Third Order
OLSR	Optimized Link State Routing
OSI	Open System Interconnection
OSPF	Open Shortest Path First
OWL	Ontological Web Language
PAR	Peak to Average Ratio

PCR	Policy Conformance Reasoner
PD	Probability of Detection
PDF	Probability Density Function
PER	Packet Error Rate
PFA	Probability of False Alarm
PHY	Physical Layer
PLNFF	Pick Lowest Noise Floor First
PQBF	Pick Quietest Band First
PSAT	Power, Saturated
PSK	Phase Shift Keying
Q	Quality Factor (generally used to describe bandwidth of filters
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase Shift Keying
RF	Radio Frequency
RFC	Request for Comment
RFIC	Radio Frequency Integrated Circuit
RKRL	Radio Knowledge Representation Language
ROC	Receiver Operating Curve
RSSI	Received Signal Strength Indication
RTS	Request to Send
SCC	Standards Coordinating Committee
SD	Spectrum Usage Density
SDR	Software Defined Radio
SFDR	Spur Free Dynamic Range
SNR	Signal to Noise Ratio
SINR	Signal to Interference and Noise Ratio
SPD	Spectral Power Density
SSI	Single Source Interference
SSID	Service Set Identifier (802.11)
SSR	System Strategy Reasoner
TCP	Transmission Control Protocol
TDMA	Time Division Multiple Access
TV	Television
UDP	User Datagram Protocol
UHF	Ultra High Frequency
URL	Uniform Resource Location
US	United States of America
UWB	Ultra Wideband
VHF	Very High Frequency
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network
WAN	Wide Area Network

Wi-Fi	Wireless Fidelity (IEEE 802.11a, b, g, n, etc.)
WiMax	Worldwide Interoperability for Microwave Access
WISP	Wireless Internet Service Provider
WLAN	Wireless Local Area Network
WNaN	Wireless Network after Next Program (DARPA)
XG	Next Generation Communications Program (DARPA)
XML	Extensible Markup Language

Appendix B Equation Variables

Symbol	**Name**	**Definition**
A_{Manual}	Manual Deconfliction Area	Area that is deconflicted for individual radios operating in manual spectrum assignment.
A_{Mobility}	Operating Region	Two dimensional area over which nodes operate. In same units as R_{wc}.
A_{spectrum}	Spectrum Availability	Probability of spectrum availability.
a_{t}	Amplitude Threshold	Threshold signal density for declaring a channel as occupied. Typically dBm/Hertz for the resolution bandwidth of the sensor.
B	Bandwidth	Bandwidth for Shannon channel capacity .
b_0	Occupied Bandwidth (in MHz)	Signal bandwidth which will be the basis for signal to noise determination.
$bandusage_i$	User Bandwidth	The instantaneous spectrum usage used by user i.
$bandwidth_0$	Available Spectrum Bandwidth	The total bandwidth made available to a set of users.
$Benefit_{\mathrm{FENoise}}$	Benefit of Cognitive Radio Noise Reduction	The improvements in overload Intermodulation through cognitive adaptations.
$Benefit_{\mathrm{Overload}}$	Benefit of Cognitive Radio overload reduction	The improvements in overload probability through cognitive adaptations.
BW	Pre-selector Bandwidth Ratio	Pre-selector bandwidth that will be provided to the first (or later) receiver amplifier stages, expressed as a proportion of the center frequency of the filter.
C	Channel Capacity	Capacity for Shannon-bound channel.
$d_{\mathrm{C\text{-}C}}$	Cognitive Radio Separation	Distance between the two Cognitive Radios.
$d_{\mathrm{NC\text{-}CRR}}$	Cognitive and non-Cognitive Radio Separation	Distance between the Cognitive Radio receiver and the non-cooperative transmitter.
$d_{\mathrm{T\text{-}T}}$	Transmitter Separation	Distance between the two Transmitting Radios.
$duty$	Network Duty Cycle	Ratio of time transmitting on a given network to the total time (aggregated across nodes).
Eff_{PA}	PA Efficiency	Efficiency of Power Amplifier.
f_{c}	Filter Center Frequency	The middle of the filter range $fc = f_{\mathrm{low}} + (f_{\mathrm{high}} - f_{\mathrm{low}})/2$.
FE_{max}	Maximum Front-end Pre-selector Power	The maximum power within the specified BW parameter between f_{low} and f_{high}.
FE_{min}	Minimum Front-end Pre-selector Power	The minimum power within the specified BW parameter between f_{low} and f_{high}.
FE_{powerin}	Front-end Input Power	The total power provided to the LNA and later stages.

Symbol	Name	Definition
$FE\alpha$	Front-end Energy Distribution alpha parameter	The Beta distribution alpha (α) parameter for a given pre-selector bandwidth parameter.
$FE\beta$	Front-end Energy Distribution beta parameter	The Beta distribution beta (β) parameter for a given pre-selector bandwidth parameter.
f_{high}	Frequency – Highest	The highest frequency of the highest pre-selector setting.
f_{low}	Frequency –Lowest	The lowest frequency of the lowest pre-selector setting.
F_{margin}	Fade Margin	Degree of additional fade margin added to the link budget in the worst-case range condition. 1 = no fading, 10^{-2} = 20 dB margin.
G_{LNA}	LNA Gain	The Gain of the LNA Stage.
I	Individual Node Throughput	The mean bandwidth available (can be sourced) from a node to a random node within a network. Bandwidth used for relayed or routed traffic is not included in this value.
I_{Density}	Spectrum Density Index	Arbitrary index of the density of signals within a specified spectrum.
I_{DSA}	DSA Performance Index	The ratio of rendezvous time to sensing time.
$I_{\mathrm{Intensity}}$	Spectrum Intensity Index	Arbitrary index of the intensity of signals within a specified spectrum.
$IIP2$	2nd Order Intercept	The effective 2nd Order Intercept of all stages in front of the modem.
$IIP3$	3rd Order Intercept	The effective 3rd Order Intercept of all stages in front of the modem.
k_{FEmaxi}	Front-end Maximum Power Coefficient	The *i*th order bandwidth polynomial coefficient for the maximum energy of the Front-end Energy Beta distribution.
$k_{\mathrm{FEmin}i}$	Front-end Minimum Power Coefficient	The *i*th order bandwidth polynomial coefficient for the minimum energy of the Front-end Energy Beta distribution.
$k_{\mathrm{FE}\alpha i}$	Front-end alpha Coefficient	The *i*th order bandwidth polynomial coefficient for the alpha term of the Front-end Energy Beta distribution.
$k_{\mathrm{FE}\beta i}$	Front-end beta Coefficient	The *i*th order bandwidth polynomial coefficient for the Beta term of the Front-end Energy Beta distribution.
K_{path}	Path Loss Constant	Path loss during first unit distance propagation.
k_{SOmaxi}	Channel Maximum Power Coefficient	The *i*th order bandwidth polynomial coefficient for the maximum energy of the Channel Occupancy Beta distribution.
$k_{\mathrm{SOmin}i}$	Channel Minimum Power Coefficient	The *i*th order bandwidth polynomial coefficient for the minimum energy of the Channel Occupancy Beta distribution.

Symbol	**Name**	**Definition**
$k_{\mathrm{SO}\alpha i}$	Channel alpha Coefficient	The *i*th order bandwidth polynomial coefficient for the alpha term of the Channel Occupancy Beta distribution.
$k_{\mathrm{SO}\beta i}$	Channel beta Coefficient	The *i*th order bandwidth polynomial coefficient for the Beta term of the Channel Occupancy Beta distribution.
L_{margin}	Link Margin	Ratio of energy available at the receiver in excess of that required to demodulate to an acceptable BER to account for all link disturbances. Includes fade margin.
n	Number of Nodes	The quantity of nodes distributed within a given region, or as members of a network.
N_0	Noise Floor	Total receiver generated and environmental noise.
Needed	Frequency Assignments Needed	Number of frequency assignments that are needed by a pool of nodes.
N_{Pool}	Pool Size	Number of channels that are contended for among the spectrum users.
$N_{\mathrm{Recipients}}$	Number of Units	Number of Units that must receive a broadcast.
P_{bo}	IMD3 Noise within b_0 bandwidth	The total mean power distributed across a signaling bandwidth b_0.
P_{D}	Probability of Detection	The probability of a single sensor scan detecting energy in a channel.
P_{FA}	Probability of False Alarm	The probability of a single sensor scan falsely detecting energy in a channel.
$P_{FE\mathrm{OverloadCR}}$	Non-Cognitive Radio Probability of Front-end Overload	Probability that a non-cognitive radio front-end will be overloaded in a given environment and with stated IIP3 and BW parameters.
$P_{\mathrm{FEOverloadNCR}}$	Cognitive Radio Probability of Front-end Overload	Probability that a cognitive radio front-end will be overloaded in a given environment and with stated IIP3 and BW parameters.
P_{IMD3}	IMD3 Noise Power	Total IIP3 generated noise, expressed as the equivalent signal input power.
$P_{\mathrm{interference}}$	Interference Threshold	Maximum energy present on a channel in order for that channel to be reused.
$P_{\mathrm{interruption}}$	Probability of Interruption	Probability that the link will fail for environmentally induced causes.
P_{NC}	Power of a Non-Cooperative Transmitter	Effective isotropic transmit power of a non-cooperative node.
P_{PA}	Power Amplifier Out	Power Amplifier output power.
$P_{\mathrm{rcvanalog}}$	Analog Receiver section power	Power for the analog portion of a receiver.
P_{rcvdig}	Digital Section Receiver Power	Power consumption of the Digital Usage in Receiver (A to D and beyond).

Symbol	**Name**	**Definition**
P_{recieve}	Required Receive Energy	Required energy at the receiver in the worst case range condition.
P_{signal}	Signal Power	Signal, or in-band power.
PSS	Pres-Selector settings	The number of statistically independent pre-selector settings that are available within tuning range, and containing candidate frequencies of operation.
P_{T}	Transmitted Power	Transmitted Power from Antenna.
P_{xmit}	Transmit Power Consumption	Transmitter section power consumption, not including PA.
R_{wc}	Worst Case Range	Required range corresponding to the maximum fade and the worst case α.
S	Signal Energy	Signal energy for Shannon capacity determination.
S_{DSA}	DSA Provided Node Separation Distance	Separation required for nodes operating under DSA principles.
SIE	Spectrum Information Effectiveness	The overall effectiveness of spectrum usage measured proportional to bits/Hertz/Area.
$SINR$	Signal to Interference and Noise Ratio	Ration of Signal Energy to the sum of Interference and Noise.
S_{manual}	Manually Planned Node Separation distance	Separation required for nodes operating under manually deconflicted spectrum principles.
$SO_{\max}$	Maximum Channel Occupancy Energy	The maximum power within any bandwidth in a range of frequencies and location.
$SO_{\min}$	Minimum Channel Occupancy Energy	The minimum power within any bandwidth in a range of frequencies and location.
SOP	Spectrum Outage Probability	The probability that spectrum energy density will exceed a fixed threshold.
S_{RCR}	Signal at Receiving Cognitive Radio	The signal energy perceived at a receiving cognitive radio.
S_{TCR}	Signal at Transmitting Cognitive Radio	The signal energy perceived at a transmitting cognitive radio.
t_{reform}	Network Reformation Time	Time, after transmissions terminate, for the Dynamic Spectrum Device to become available for application traffic.
t_{release}	Abandonment Time	Time, following t_{sense}, for the Dynamic Spectrum Device to abandon the channel, measured when transmissions terminate.
$t_{\mathrm{rendezvous}}$	Rendezvous Time	The mean time for a DSA system to abandon a frequency, and move the network to a new frequency.
t_{sense}	Sensing Time	Time interval between the start of Dynamic Spectrum Devices sensing periods, or the time to complete a cycle of all monitored frequencies.

Symbol	**Name**	**Definition**
t_{sensing}	Sensing Interval	Time to sense the spectrum to detect signals within the band of interest.
v_{in}	Input Voltage	Signal voltage presented to the input of the LNA.
v_{out}	Output Voltage	Amplified signal voltage.
$xmit$	Transmit Duty Cycle	Duty Cycle of the Transmitter Section of a radio.
α	Propagation Exponent	Reflecting the path loss between transmitter and receiver. Typically between 2 and 3.8.
α_{bc}	Best Case Propagation[1]	Propagation coefficient corresponding to the best case conditions.
α_{com}	Communication Link Propagation Exponent	The propagation exponent between link partners, i.e., the transmitter and the intended receiver(s).
α_{int}	Interference Propagation Exponent	The propagation exponent between a link transmitter and the unintended receiver(s), for which this emission constitutes interference.
α_{mean}	Mean Propagation Exponent	The mean value of the propagation exponent within a given set of links.
α_{minimum}	Minimum Propagation Exponent	The minimum of the propagation exponent within a given set of links.
α_{tc}	Typical Propagation	Propagation coefficient corresponding to the typical conditions.
α_{variance}	Propagation Exponent Variance	The variance of the propagation exponent within a given set of links.
α_{wc}	Worst Case Propagation	Propagation coefficient corresponding to the worst-case conditions.
$\mathcal{P}(\,)$	Probability Distribution	A Probability distribution of the parametric variable.

---

[1] The units of distance for propagation are arbitrary, but should be selected such that they are similar to the distance at which the best and worst-case propagation diverge; i.e., the point at which propagation transitions from direct to diffraction. Typically, this might be 200 to 500 meters for a low antenna, 1 km for a handheld, and would be irrelevant for a space link, since all propagation would be r^2 at reasonable "look angles" or elevation from the surface.

References

[1] W. W. Massie and C. R. Underhill, "The Future of the Wireless Art," *Wireless Telegraphy and Telephony*, pp. 67–71, 1908.

[2] L. Kleinrock, *Communication Nets: Stochastic Message Flow and Delay*. New York: McGraw-Hill, 1964.

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

[4] J. Mitola and G. Maguire, "Cognitive Radio: Making Software Radios More Personal," in *IEEE Personal Communications*, vol. 6, no. 4, pp. 13-18, 1999.

[5] J. Mitola, "Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio," Royal Institute of Technology (KTH), Sweden, Dissertation, Doctor of Technology, May 8, 2000.

[6] P. Marshall, "Extending the Reach of Cognitive Radio," *Proceedings of the IEEE*, vol. 97, no. 4, pp. 612-625, Apr. 2009.

[7] P. Marshall, "Cognitive Radio as a Mechanism to Manage Front-End Linearity and Dynamic Range," *IEEE Communications Magazine*, vol. 47, no. 3, pp. 81-87, Mar. 2009.

[8] P. S. Ryan, "Wireless Communications and Computing at a Crossroads: New Paradigms and Their Impact on Theories Governing the Public's Right to Spectrum Access," *Journal on Telecommunications & High Technology Law*, vol. 3, No. 2, p. 239, 2005.

[9] R. Sofia and P. Mendes, "User-Provided Networks: Consumer as Provider," *IEEE Communications*, vol. 46, no. 12, pp. 86-91, Dec. 2008.

[10] Defense Advanced Research Projects Agency, Wireless Network after Next, Proposer Information Pamphlet, http://www.darpa.mil/sto/solicitations/WNaN/pip.htm.

[11] P. Marshall, "DARPA Progress towards Affordable, Dense, and Content Focused Tactical Edge Networks," in *IEEE Military Communications Conference*, San Diego, CA, Nov. 2008.

[12] P. Marshall, "What's Beyond Software Defined Radios?," in *2003 Software Defined Radio Forum Annual Technical Conference*, Orlando, FL, 2003.

[13] P. Marshall, "Cognitive Radio: What is the Value Equation?," in *2004 Software Defined Radio Forum Annual Technical Conference*, Phoenix, AZ, Nov. 18, 2004.

[14] E. Leitl, "Information Technology Issues During and after Katrina and Usefulness of the Internet: How We Mobilized and Utilized Digital Communications Systems," *Critical Care*, vol. 10(1):110, 2006.

[15] M. Donahoo and B. Steckler, "Emergency mobile wireless networks," in *IEEE Military Communications Conference*, pp. 2413-2420, Oct. 17-20, 2005.

[16] J. Neel, "Analysis and Design of Cognitive Radio Networks and Distributed Resource Management Algorithms," Virginia Polytechnic Institute, Blacksburg, VA, Ph.D Dissertation 2006.

[17] QinetiQ, "Cognitive Radio Technology: A Study for OFCOM – Summary Report," QINETIQ/06/00420, Feb. 2007.

[18] F. Weidling, D. Datla, V. Petty, P. Krishnan, and G. Minden, "A Framework for R.F. Spectrum Measurements and Analysis," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, 2005.

[19] IEEE, "IEEE Standard Definitions and Concepts for Dynamic Spectrum Access: Terminology Relating to Emerging Wireless Networks, System Functionality, and Spectrum Management," IEEE Std 1900.1™-2008, Sept. 26, 2008.

[20] S. Haykin, "Fundamental Issues in Cognitive Radio," in *Cognitive Radio Networks*. Springer-Verlag, 2007.

[21] Q. Zhao and B. Sadler, "A Survey of Dynamic Spectrum Access," *IEEE Signal Processing Magazine*, May 2007.

[22] I. Akyildiz, W. Lee, M. Vuran, and S. Mohanty, "NeXt Generation Dynamic Spectrum Access Cognitive Radio Wireless Networks: A Survey," *IEEE Communications Magazine*, vol. 40-48, pp. 40-48, Apr. 2008.

[23] A. Ozgur, O. Leveque, and D. Tse, "How does Information Capacity of Ad-Hoc Networks Scale?," in *44th Annual Allerton Conference on Communication, Control and Computing*, 2006.

[24] A. Pereira, J. Bastos, P. Marques, and A. Gameiro, "Analysis of Communication Opportunities in UMTS Cellular Networks," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 99-102, Apr. 17-20, 2007.

[25] C-T Chou, S. Shankar, H. Kim, and K. Shin, "What and How Much to Gain by Spectrum Agility?," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 3, Apr. 2007.

[26] D. Sicker, "The Technology of Dynamic Spectrum Access and its Challenges," *IEEE Wireless Magazine*, vol. 45, no. 6, June 2007.

[27] J. Neel and J. Reed, "Performance of Distributed Dynamic Frequency Selection Schemes for Interference Reducing Networks," in *IEEE Military Communications Conference*, 2006.

[28] IEEE 802.22 Working Group on Wireless Regional Area Networks, *Functional Requirements for the 802.22 WRAN Standard*, http:www.ieee802.org/22.

[29] C. Cordeiro, K. Challapali, B. Dagnachew, and S. Shankar, "IEEE 802.22: An Introduction to the First Wireless Standard based on Cognitive Radios," *IEEE Journal of Communications*, vol. 1, no. 1, pp. 38-47, Apr. 2007.

[30] C. Cordiero, K. Challapali, D. Birru, and S. Shankar, "IEEE 802.22: The First World Wide Standard for Cognitive Radio," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, 2005.

[31] T. Rondeau, C. W. Bostian, D. Maldonado, A. Ferguson, S. Ball, and B. Le, "Cognitive Radios in Public Safety and Spectrum Management," in *Telecommunications Policy and Research Conference*, Sept. 2005.

[32] M. Silvius, F. Ge, A. Young, A. MacKenzie, and C. Bostian, "Smart Radio: Spectrum Access for First Responders," in *Proceedings of SPIE - The International Society for Optical Engineering*, 2008.

[33] S. Jones, E. Jung, L. Xin, N. Merheb, and I-Jeng Wang, "Characterization of Spectrum Activities in the US Public Safety Band for Opportunistic Spectrum Access," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Apr. 17-20, 2007.

[34] S. Chan, "Shared Spectrum Access for the DoD," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Apr. 17-20, 2007.

[35] P. Amini, E. Azarnasabi, S. Akoum, and B. Farhang-Boroujeny, "An Experimental Cognitive Radio for First Responders," in *3rd IEEE Intermational Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Oct. 14-17, 2008.

[36] M. Buddhikot and K. Ryan, "Spectrum Management in Coordinated Dynamic Spectrum Access Based Cellular Networks," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 299-307, Nov. 8-11, 2005.

[37] A. Subramanian, M. Al-Ayyoub, H. Gupta, S. Das, and M. Buddhikot, "Near-Optimal Dynamic Spectrum Allocation in Cellular Networks," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Oct. 2008.

[38] United States Federal Communications Commission (FCC), "Notice for Proposed Rulemaking (NPRM 03 322): Facilitating Opportunities for Flexible, Efficient and Reliable Spectrum Use Employing Spectrum Agile Radio Technologies," ET Docket No. 03 108, Dec. 2003.

[39] Federal Communications Commission, "In the Matter of Unlicensed Operation in the TV Broadcast Bands Additional Spectrum for Unlicensed Devices Below 900 MHz and in the 3 GHz Band," Nov. 14, 2008.

[40] M. McHenry, K. Steadman, and M. Lofquist, "Determination of Detection Thresholds to Allow Safe Operation of Television Band “White Space” Devices," in *3rd IEEE International Symposium on Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.

[41] J. Park, T. Song, J. Hur, S. M. Lee, J. Choi, K. Kim, J. Lee, K. Lim, C. H. Lee, H. Kim, and J. Lasker, "A Fully Integrated UHF Receiver with Multi-resolution Spectrum Sensing (MRSS) Functionality for IEEE 802.22 Cognitive Radio Applications," in *2008 IEEE International Solid State Circuits Conference*, Feb. 2008.

[42] Y. Hur, J. Park, W. Woo, J. S. Lee, K. Lim, C. H. Lee, H. S. Kim, and J. Laskar, "A Cognitive Radio System Employing A Dual-Stage Spectrum Sensing Technique: A Multi-Resolution Spectrum Sensing and A Temporal Signature Detection Technique," in *Global Telecommunications Conference*, pp. 1-5, Nov. 2006.

[43] F. Ge and C. Bostian, "A Parallel Computing Based Spectrum Sensing Approach for Signal Detection under Conditions of Low SNR and Rayleigh Multipath Fading," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago IL, Oct. 2008.

[44] W. Hu, D. Willkomm, M. Abusubaih, J. Gross, G. Vlantis, M. Gerla, and A. Wolisz, "Dynamic Frequency Hopping for Efficient IEEE 802.22 Operation," in *IEEE Communications Magazine*, May 2007.

[45] V. Petty, R. Rajbanshi, D. Datla, F. Weidling, D. DePardo, P. Kolodzy, M. Marcus, A. M. Wyglinski, J. Evans, G. Minden, and J. A. Roberts, "Feasibility of Dynamic Spectrum Access in Underutilized Television Bands," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 331-339, Apr. 2007.

[46] T. X. Brown and A. Sethi, "Potential Cognitive Radio Denial-of-Service Vulnerabilities and Countermeasures," *Proceedings of the International Symposium on Advanced Radio Technologies*, Boulder, CO, pp. 44-51, Feb. 2007.

[47] R. Chen, J. Park, Y. Hou, and J. Reed, "Toward Secure Distributed Spectrum Sensing in Cognitive Radio Networks," *IEEE Communications Magazine*, pp. 50-55, Apr. 2008.

[48] C. Rieser, T. Rondeau, C. Bostian, and T. Gallagher, "Cognitive Radio Testbed: Further Details and Testing of a Distributed Genetic Algorithm Based Cognitive Engine for Programmable Radios," in *Proceedings of the IEEE Military Communications Conference*, 2004.

[49] T. W. Rondeau and C. W. Bostian, "Cognitive Techniques: Physical and Link Layers," in *Cognitive Radio*. B. Fette, Ed. Elsevier/Newnes, 2006.

[50] T. Rondeau, B. Le, C. Rieser, and C. Bostian, "Cognitive Radios with Genetic Algorithms: Intelligent Control of Software Defined Radios," in *Software Defined Radio Forum Annual Conference*, 2004.

[51] T. Rondeau, "Application of Artificial Intelligence to Wireless Communications," Virginia Polytechnic Institute and State University, Ph.D Dissertation, Sept. 20, 2007.

[52] A. Mody, S. Blatt, N. Thammakhoune, T. McElwain, J. Niedzwiecki, D. Mills, M. Sherman, and C. Myers, "Machine Learning Based Cognitive Communications in White as well as the Gray Space," in *IEEE Military Communications Conference*, vol. 26, no. 1, pp. 2321-2327, Oct. 2007.

[53] A. N. Mody, S. R. Blatt, D. G. Mills, T. P. McElwain, N. B. Thammakhoune, J. D. Niedzwiecki, M. J. Sherman, C. S. Myers, and P. D. Fior, "Recent Advances in Cognitive Communications," *IEEE Communications Magazine*, Oct. 2007.

[54] Z. Zhang and X. Xie, "Application research of evolution in cognitive radio based on GA," in *3rd IEEE Conference on Industrial Electronics and Applications*, pp. 1575-1579, June 3-5, 2008.

[55] C. Clancy, J. Hecker, E. Stuntebeck, and T. O'Shea, "Applications of Machine Learning to Cognitive Radio Networks," *IEEE Wireless Communications*, vol. 14, no. 4, Aug. 2007.

[56] F. Ge, Q. Chen, Y. Wang, C. Bostian, T. Rondeau, and B. Le, "Cognitive Radio: From Spectrum Sharing to Adaptive Learning and Reconfiguration," in *IEEE Aerospace Conference*, 2008.

[57] T. Rondeau, B. Le, D. Maldonado, D. Scaperoth, and C. Bostian, "Cognitive Radio Formulation and Implementation," in *Cognitive Radio Oriented Wireless Networks and Communications*, June 2006.

[58] Y. Su and M. van der Schaar, "Learning for Cognitive Wireless Users," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Oct. 14-17, 2008.

[59] S. Yarkin and H. Arslan, "Exploiting Location Awareness toward Improved Wireless Systems Design in Cognitive Radio," *IEEE Communications Magazine*, vol. 46, no. 1, pp. 128-135, Jan. 2008.

[60] K. Nolan, P. Sutton, and L. Doyle, "An Encapsulation for Reasoning, Learning, Knowledge Representation, and Reconfiguration Cognitive Radio Elements," in *1st International Conference on Cognitive Radio Oriented Wireless Networks and Communications*, pp. 1-5, June 8-10, 2006.

[61] R. Thomas, L. DaSilva, and A. MacKenzie, "Cognitive Networks," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, 2005.

[62] R. Thomas, D. Friend, L. Dasilva, and A. Mackenzie, "Cognitive Networks: Adaptation and Learning to Achieve End-to-End Performance Objectives," *IEEE Communications Magazine*, vol. 44, no. 12, Dec. 2006.

[63] IEEE Standards Coordinating Committee 41, "Dynamic Spectrum Access Networks," 2007, http://www.scc41.org/.

[64] M. Sherman, A. Mody, R. Martinez, C. Rodrigues, and R. Reddy, "IEEE Standards Supporting Cognitive Radio and Networks, Dynamic Spectrum Access, and Coexistence," *IEEE Communications Magazine*, vol. 46, no. 7, pp. 72-79, July 2008.

[65] W. Krenik, A. M. Wyglinski, and L. E. Doyle, "Cognitive Radio for Dynamic Spectrum Access," *IEEE Communications Magazine*, vol. 45, no. 5, May 2007.

[66] K. E. Nolan, P. D. Sutton, L. E. Doyle, T. W. Rondeau, B. Le, and C. W. Bostian, "Demonstration and Analyses of Collaboration, Coexistence, and Interoperability of Cognitive Radio Platforms," in *4th IEEE Consumer Communications and Networking Conference*, Jan. 11-13, 2007.

[67] D. Raychaudhuri, N. Mandayam, J. Evans, B. Ewy, S. Seshan, P. Steenkiste, "CogNet – An Architectural Foundation for Experimental Cognitive Radio Networks within the Future Internet," in *First International Workshop on Mobility in the Evolving Internet Architecture*, 2006.

[68] Y. Yuan, P. Bahl, R. Chandra, T. Moscibroda, and Y. Wu, "Allocating Dynamic Time-Spectrum Blocks in Cognitive Radio Networks," in *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Montreal, Canada, Sept. 2007.

[69] J. Zhao, H. Zheng, and G. Yang, "Distributed Coordination in Dynamic Spectrum Allocation Networks," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005.

[70] D. Scaperoth, B. Le, T. Rondeau, D. Maldonado, C. Bostian, and S. Harrison, "Cognitive Radio Platform Development for Interoperability," in *IEEE Military Communications Conference*, pp. 2640-2645, Oct. 2006.

[71] P. Marshall, "XG Program Field Test Results," in *Defense Spectrum Summit 2007*, Nov. 2007.

[72] United States Federal Communications Commission, "Facilitating Opportunities for Flexible, Efficient, and Reliable Spectrum Use Employing Cognitive Radio Technologies," ET Docket No. 03-108, Dec. 17, 2003.

[73] Office of Communications (OFCOM) (UK), "A Study into Dynamic Spectrum Access," London, UK, Final Report, Report No. 72/06/R/353/U, Mar. 7, 2007.

[74] D. Elenius, G. Denker, M. Stehr, R. Senanayake, C. Talcott, and D. Wilkins, "CoRaL - Policy Language and Reasoning Techniques for Spectrum Policies," in *IEEE Workshop on Policies for Distributed Systems and Networks*, June 2007.

[75] F. Perich, "Policy-based Network Management for NeXt Generation Spectrum Access Control," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Apr. 2007.

[76] D. Wilkins, G. Denker, M. Stehr, D. Elenius, R. Senanayake, and C. Talcott, "Policy-Based Cognitive Radios," *IEEE Wireless Communications*, vol. 14, no. 4, pp. 41-46, Aug. 2007.

[77] P. Marshall, "XG Communications Program Information Briefing," in *Semantic Web Applications for National Security (SWANS) Conference*, Apr. 2005.

[78] M. McHenry, E. Livsics, T. Nguyen, and N. Majumdar, "XG Dynamic Spectrum Access Field Test Results," *IEEE Communications Magazine*, vol. 45, no. 6, pp. 51–57, June 2007.

[79] S. Doshi, D. Sheetalkumar, H. Duong, R. Bagrodia, and S. Thai, "Spectrum and Network Management Convergence for Wireless Communications," in *Military Communications Conference*, Washington, DC, Oct. 2006.

[80] F. Cohen, "Policy-Based Security and Enterprise Policy Management," Burton Group, Research Report, Dec. 2003.

[81] D. Verma, *Policy-Based Networking: Architecture and Algorithms*. New Riders Publishing, 2001.

[82] J. Strassner, *Policy-Based Network Management: Solutions for the Next Generation*. Morgan Kaufmann, 2004.

[83] P. Marshall, "DARPA Next Generation (XG) Communications Program," in *Forum on Spectrum Management Policy Reform, National Academy of Engineering, Computer Science and Telecommunications Board (CSTB)*, Washington, DC, Feb. 13, 2004.

[84] T. Issariyakul, L. Pillutla, and V. Krishnamurthy, "Tuning Radio Resource in an Overlay Cogntive Radio Nework for TCP: Greed Isn't Good," *IEEE Communications Magazine*, vol. 47, no. 7, pp. 57-63, July 2009.

[85] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized Cognitive MAC for Opportunistic Spectrum Access in Ad Hoc Networks: a POMDP Framework," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 3, pp. 589-600, Apr. 2007.

[86] Y. Zhao, S. Mao, J. Neel, and J. Reed, "Performance Evaluation of Cognitive Radios: Metrics, Utility Functions and Methodologies," *Proceedings of the IEEE – Special Issue on Cognitive Radio*.

[87] P. Marshall, "Recent Progress in Moving Cognitive Radio and Services to Deployment," in *9th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, Newport Beach, CA, June 2008.

[88] P. Marshall, "Protocols for Spectrum Sharing," in *2002 Software Defined Radio Forum Annual Technical Conference*, San Diego, CA, Nov. 12, 2002.

[89] P. Marshall, "Policy and Autonomous Behavior within Wireless Devices," in *Federal Communications Commission (FCC) Technical Advisory Committee (TAC)*, Washington, DC, July 28, 2005.

[90] G. Denker, D. Elenius, R. Senanayake, M-O. Stehr, and D. Wilkins, "A Policy Engine for Spectrum Sharing," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 55-65, Dublin, IE, 2007.

[91] M. Kokar, "The Role of Ontologies in Cognitive Radio," in *Cognitive Radio*, B. Fette, Ed., Elsevier, 2006.

[92] United States Federal Communications Commission, "Revision of Part 15 of the Commission's Rules Regarding Ultra-Wideband Transmission Systems," ET Docket No. 98-153, Released Dec. 16, 2004.

[93] Z. Wu and B. Natarajan, "Interference Tolerant Agile Cognitive Radio: Maximize Channel Capacity of Cognitive Radio," in *4th IEEE Consumer Communications and Networking Conference*, Jan. 2007.

[94] Z. Wu, P. Ratazzi, V. Chakravarthym, and L. Hong, "Performance Evaluation of Adaptive Non-contiguous MC-CDMA and Non-contiguous CI/MC-CDMA for Dynamic Spectrum Access," in *IEEE Cognitive Radio Oriented Wireless Networks and Communications Conference*, Singapore, 2008.

[95] V. Chakravarthy, M. A. Temple, and Z. Wu, "A General Overlay/Underlay Analytic Expression Representing Cognitive Radio Waveforms," in *International Waveform Diversity & Design Conference*, Italy, 2007.

[96] R. Rajbanshi, Q. Chen, A. Wyglinski, G. Milden, and J. Evans, "Quantitative Comparision of Agile Modulation Technique for Cognitive Radio Tranceivers," in *4th IEEE Consumer Communications and Networking Conference*, Jan. 2007.

[97] S. Cui, A. J. Goldsmith, and A. Bahai, "Energy-Constrained Modulation Optimization," *IEEE Transactions on Wireless Communications,* vol. 4, no. 5, pp. 2349-2360, Sept. 2005.

[98] P. Marshall, "Wireless Networking Impacts on Network Protocols," in *IEEE Computer Network Computing and Applications Symposium*, Cambridge, MA, Aug. 31, 2004.

[99] P. Marshall, "Policy-Based Cognitive Radios: Future Framework for Spectrum Access," in *United Nations, International Telecommunications Union (ITU-R), Radio Communications Bureau Seminar on Tomorrow's Technological Innovations Related to the World Radio Conference 2007*, Geneva, Switzerland, Sept. 9, 2004.

[100] G. D. Golden, G. J. Foschini, R. A. Valenzuela, and P. W. Wolniansky, "V-BLAST: A High Capacity Space-Time Architecture for the Rich-Scattering Wireless Channel," in *Fifth Workshop on Smart Antennas in Wireless Mobile Communications*, Stanford, CA, July 23-24, 1998.

[101] P. Marshall, "Dynamic Spectrum Management of Front-End Linearity and Dynamic Range," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, 2008.

[102] Z. Ji and K. Liu, "Cognitive Radios for Dynamic Spectrum Access - Dynamic Spectrum Sharing: A Game Theoretical Overview," *IEEE Communications Magazine*, vol. 45, no. 5, pp. 88-94, May 2007.

[103] S. Fischer, P. Mähönen, R. Fischer, P. Mahonen, M. Schoengens, and B. Vöcking, "Load Balancing Dynamic Spectrum Assignment with Local Information for Secondary Users," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.

[104] P. Marshall, "From Self-Forming Mobile Networks to Self-Forming Content Networks," in *Association of Computing Machinery Mobile Communications Conference*, San Francisco, CA, Sept. 2008.

[105] S. Govindasamy, D. Bliss, and D. Staelin, "Spectral Efficiency in Single-Hop, Ad-Hoc Wireless Networks with Interference Using Adaptive Antenna Arrays," *IEEE Journal on Selected Areas in Communications*, 2007.

[106] A. Goldsmith, S. A. Jafar, I. Maris, and S. Srinivasa, "Breaking Spectrum Gridlock with Cognitive Radios: an Information Theoretic Perspective," *Proceedings of the IEEE*, 2008.

[107] P. Marshall, "Adaptation and Integration Across the Layers of Self Organizing Wireless Networks," in *2006 Military Communications (MILCOM) Conference.* Armed Forces Communications and Electronics Association (AFCEA*)*, Washington, DC, Oct. 23, 2006.

[108] S. Oh, M. Gerla, and Joon-Sang Park, "MIMO and TCP: A Case for Cross Layer Design," in *IEEE Military Communications Conference*, Orlando, FL, Oct. 29-31, 2007.

[109] D. Perkins and H. Hughes, "TCP Performance in Mobile Ad Hoc Networks," in *Proceedings of the IEEE Wireless Communications and Networking Conference*, 1999.

[110] M. Gerla, K. Tang, and R. Bagrodia, "TCP Performance in Wireless Muti-Hop Networks," in *Proceedings of the IEEE Mobile Computing Systems and Applications Conference*, Feb. 1999.

[111] G. Chan and B. Brown, "A Measure of Spectrum Utilization in Land Mobile Bands," in *Canadian Conference on Electrical and Computer Engineering 1996,* vol. 1, May 1996.

[112] R. Bacchus, A. Fertner, C. Hood, and D. Roberson, "Long-term, wide-band spectral monitoring in support of Dynamic Spectrum Access Networks at the IIT Spectrum Observatory," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.

[113] M. McHenry, P. Tenhula, D. McCloskey, D. Roberson, and C. Hood, "Chicago Spectrum Occupancy Measurements & Analysis and a Long-term Studies Proposal," in *Proceedings of the First International Workshop on Technology and Policy for Accessing Spectrum*, Boston, MA, 2006.

[114] M. McHenry and K. Steadman, "Spectrum Occupancy Measurements, Location 1 of 6: Riverbend Park, Great Falls, Virginia," National Science Foundation, Shared Spectrum Company Report, Aug. 2005.

[115] M. McHenry and K. Steadman, "Spectrum Occupancy Measurements, Location 2 of 6: Tysons Square Center, Vienna, Virginia, April 9, 2004," National Science Foundation, Shared Spectrum Company Report, Aug. 2005.

[116] M. McHenry and K. Steadman, "Spectrum Occupancy Measurements, Location 4 of 6: Republican National Convention, New York City, New York, August 30, 2004 - September 3, 2004, Revision 2," National Science Foundation, Shared Spectrum Company Report, Aug. 2005.

[117] M. McHenry and K. Steadman, "Spectrum Occupancy Measurements, Location 5 of 6: National Radio Astronomy Observatory (NRAO), Green Bank, West Virginia, October 10-11, 2004, Revision 3," National Science Foundation, Shared Spectrum Company Report, Aug. 2005.

[118] M. McHenry and K. Steadman, "Spectrum Occupancy Measurements, Location 6 of 6: Shared Spectrum Building Roof, Vienna, Virginia, December 15-16, 2004," National Science Foundation, Shared Spectrum Company Report, Aug. 2005.

[119] National Radio Astronomy Observatory. Robert C. Byrd Green Bank Telescope (GBT), http://www.gb.nrao.edu/gbt/.

[120] M. McHenry and K. Steadman, "Spectrum Occupancy Measurements, Location 3 of 6: National Science Foundation Building Roof," National Science Foundation, Shared Spectrum Company Report, Aug. 15, 2005.

[121] A. Sahai, N. Hoven, and R. Tandra, "Some fundamental limits on cognitive radios," in *Forty-second Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2004.

[122] R. Tandra and A. Sahai, "SNR Walls for Feature Detectors," in *Proceedings of the 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, pp. 559-570, Apr. 2007.

[123] M. Abromowitz and I. A. Stegun, *Handbook of Mathematical Functions*. New York: Dover, 1970.

[124] M. Skolnik, *Introduction to Radar Systems*. 2nd ed., McGraw-Hill, Ed., 1980.

[125] Q. Zhao and A. Swami, "A Decision-Theoretical Framework for Opportunitstic Spectrum Access," *IEEE Wireless Communications*, pp. 14-20, Aug. 2007.

[126] S. Geirhofer, L. Tong, and B. Sadler, "Dynamic Spectrum Access in the Time Domain: Modeling and Exploiting White Space," *IEEE Communications Magazine*, vol. 45, no. 5, pp. 66-72, May 2007.

[127] L. Luna, "Nextel Interference Debate Rages On," in *Mobile Radio Technology*, Aug. 1, 2003.

[128] S. Bingxue, "Challenges in RF Analog Integrated Circuits," in *Proceedings of the 4th International Conference on ASIC*, pp. 800-805, 2001.

[129] R. Hickling, O. Panfilov, T. Turgeon, and M. Yagi, "Design Considerations for the Front End of Frequency Agile SDR Receivers," in *Software Defined Radio Forum Annual Technical Conference*.

[130] T.A. McKee, "Receiver intermodulation: Enforcing the square law," in *IEEE Transactions on Vehicular Technology*, vol. 18, no. 3, pp. 141-145, Nov. 1969.

[131] I. Hickman, *Practical Radio-Frequency Handbook*. Newnes, p. 161, 2007.

[132] F. Belveze and P. Baudin, "Specifying Receiver IP2 and IP3 Based on Tolerance to Modulated Blockers," *IEEE Transactions on Communications*, vol. 56, no. 10, pp. 1677-1682, Oct. 2008.

[133] Y. S. Poberezhskiy, "On Dynamic Range of Digital Receivers," in *IEEE Aerospace Conference*, 2007.

[134] H. Rhu, H. Byeon, and J. Kim, "Power Spectrum Analysis of the Cross/Intermodulation in the Non-Linear LNA," in *IEEE Transactions on Vehicular Technology*, vol. 55, no. 3, May 2006.

[135] U. L. Rhode and D. P. Newkirk, *RF/Microwave Circuit Design for Wireless Applications*. John Wiley & Sons, 2000.

[136] United States Federal Communications Commission, "Spectrum Policy Task Force Report," ET Docket No. 02-135, Nov. 2002.

[137] J. Hilbert, "RF-MEMS for Wireless Communications," *IEEE Communications Magazine*, vol. 46, no. 8, pp. 68-74, Aug. 2008.

[138] G. Rabeiz and G. L-Tan, "Introduction: RF MEMS for Microwave Applications," in *RF MEMS: Theory, Design, and Technology*. Wiley, 2003.

[139] B. Perlman, J. Laskar, and K. Lim, "Fine-Tuning Commercial and Military Radio Design," *IEEE Microwave Magazine*, vol. 9, no. 4, pp. 95-106, Aug. 2008.

[140] M. McHenry, E. Livsics, N. Thao, and N. Majumdar, "XG Dynamic Spectrum Access Field Test Results," *IEEE Communications Magazine*, vol. 45, no. 6, pp. 51-57, June 2007.

[141] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," in *IEEE Transactions on Information Theory*, vol. IT-46, no. 2, pp. 388-404, Mar. 2000.

[142] P. Marshall, "Experimental Results from Dynamic Spectrum Testing of the Defense Advanced Research Projects Agency XG Cognitive Radio System," in *2006 Software Defined Radio Forum Annual Technical Conference*, Orlando, FL, Nov. 15, 2006.

[143] L. Kovacs and A. Vidacs, "Interference-Tolerant Spatio-Temporal Dynamic Spectrum Allocation," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 403-411, Apr. 17-20, 2007.

[144] M. Win, P. Pinto, and L. Shepp, "A Mathematical Theory of Network Interference and Its Applications," *Proceedings of the IEEE*, vol. 97, no. 2, Feb. 2009.

[145] V. Marojevi, J. Salazar, K. Reves, and A. Gelonch, "On Integrating Radio, Computing, and Application Resource Management in Cognitive Radio Systems," in *Third IEEE International Conference on Wireless and Mobile Computing, Networking and Communications*, Oct. 8-10, 2007.

[146] A. Sadek, W. Zhang, and S. Shellhammer, "Listen-Before-Talk Versus Treating Interference as Noise for Spectrum Sharing," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DYSPAN)*, Chicago, IL, Oct. 2008.

[147] T. Mazraani and G. Parulkar, "Performance Analysis of the Ethernet under Conditions of Bursty Traffic," in *IEEE Global Telecommunications Conference*, pp. 592-596, 1992.

[148] A. P. Jardosh, K. N. Ramachandran, and K. Belding-Royer, "Understanding Congestion in IEEE 802.11b Wireless Networks," in *Proceedings of the Fifth ACM Internet Measurement Conference*, pp. 279–292, Berkeley, CA, Oct. 2005.

[149] D. Friend and A. MacKenzie, "Environmentally-Friendly Secondary Network Topology Control for Minimizing Outage Potential," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago IL, Oct. 2008.

[150] Y. Xing, C. Mathur, M. Haleem, R. Chandramouli, and K. Subbalakshmi, "Dynamic Spectrum Access with QoS and Interference Temperature Constraints," *IEEE Transactions on Mobile Computing*, vol. 6, no. 4, pp. 423-433, Apr. 2007.

[151] J. Bater, H. Tan, K. Brown, and L. Doyle, "Maximizing Access to a Spectrum Commons Using Interference Temperature Constraints," in *2nd International Conference on Cognitive Radio Oriented Wireless Networks and Communications*, 2007.

[152] I. Akyildiz, W. Lee, M. Vuran, and S. Mohanty, "A survey on Spectrum Management in Cognitive Radio Networks," *IEEE Communications Magazine*, vol. 46, no. 4, pp. 40-48, Apr. 2008.

[153] J. Bater, T. Hwee-Pink, K. Brown, and L. Doyle, "Modelling Interference Temperature Constraints for Spectrum Access in Cognitive Radio Networks," in *IEEE International Conference on Communications*, pp. 6493-6498, June 24-28, 2007.

[154] T. C. Clancy, "Dynamic Spectrum Access in Cognitive Radio Networks," Ph.D Dissertation, Graduate School of the University of Maryland, College Park, MD, 2006.

[155] J. B. Andersen, T. S. Rappaport, and S. Yoshida, "Propagation Measurements and Models for Wireless Communications Channels," *IEEE Communications Magazine*, vol. 33, no. 1, pp. 42-49, January 1995.

[156] A. E. Leu, M. McHenry, and B. L. Mark, "Modeling and analysis of interference in listen-before-talk spectrum access schemes," *International Journal of Network Management*, vol. 16 no 1, pp. 131–147, 2006.

[157] M. McHenry, E. Livsics, T. Nguyen, and N. Majumdar, "XG Dynamic Spectrum Sharing Field Test Results," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 17-20, 2007.

[158] F. Seelig, "A Description of the August 2006 XG Demonstrations at Fort A. P. Hill," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 17-20, 2007.

[159] M. McHenry, K. Steadman, A. Leu, and E. Melick, "XG DSA Radio System," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.

[160] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, "Delay-Tolerant Network Architecture," Internet Engineering Task Force, IETF RFC 4838, Apr. 2007.

[161] B. Le, F. Rodriguez, Q. Chen, B. Li, F. Ge, M. Nainay, T. Rondeau, and C. Bostian, "A Public Safety Cognitive Radio Node System," in *2007 Software Defined Radio Forum Technical Conference*, Nov. 2007.

[162] W. Lehr and N. Jesuale, "Spectrum Pooling for Next Generation Public Safety Radio Systems," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.

[163] W. Lehr and N. Jesuale, "Public-Safety Radios Must Pool Spectrum," *IEEE Communications Magazine*, vol. 47, no. 3, pp. 103-109, Mar. 2009.

[164] T. Weiss and F. Jondral, "Spectrum Pooling: An Innovative Strategy for the Enhancement of Spectrum Efficiency," *IEEE Communications Magazine*, Mar. 2004.

[165] N. Jesuale and B. Eydt, "A Policy Proposal to Enable Cognitive Radio for Public Safety and Industry in the Land Mobile Radio Bands," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, IE, Apr. 17-20, 2007.

[166] A. Jovicic and P. Viswanath, "Cognitive Radio: An Information-Theoretic Perspective," in *IEEE International Symposium on Information Theory*, pp. 2413-2417, July 2006.

[167] P. C. Pinto, "Communication in a Poisson Field of Interferers," Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA, Master's Thesis, 2006.

[168] P. C. Pinto and M. Z. Win, "Spectral Outage due to Aggregate Interference in a Poisson Field of Nodes," in *IEEE Global Communications Conference, 2006.*

[169] C. L. Nikias and M. Shao, *Signal Processing with Alpha-Stable Distributions and Applications*. New York: Wiley, 1995.

[170] L. DaSilva and I. Guerreiro, "Sequence-based Rendezvous for Dynamic Spectrum Access," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DYSPAN)*, Chicago, IL, Oct. 2008.

[171] C. Doerr, D. Sicker, and D. Grunwald, "What a Cognitive Radio Network could Learn from a School of Fish," in *3rd International Wireless Internet Conference*, 2007.

[172] C. Doerr, D. Grunwald, and D. Sicker, "Local Independent Control of Cognitive Radio Networks," in *3rd International Conference on Cognitive Radio Oriented Wireless Networks and Communications*, 2008.

[173] C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, pp. 379–423 and pp. 623–656, July and Oct.1948.

[174] P. Marshall, "Spectrum Awareness," in *Cognitive Radio Technology*. B. Fette, Ed., Elsevier:Boston, MA, First Edition, 2007

[175] W. Lehr and J. Crowcroft, "Managing Shared Access to a Spectrum Commons," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 420-444, Nov. 8-11, 2005.

[176] J. Crowcroft, R. Gibbens, and F. Kelly, "Modelling Incentives for Collaboration in Mobile Ad Hoc Networks," in *Proceedings of Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, Mar. 2003.

[177] M. Cooper, "The Economics of Collaborative Production in the Spectrum Commons," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 379-400, Baltimore, MD, Nov. 8-11 2005.

[178] W. Lehr, "Economic Case for Dedicated Unlicensed Spectrum Below 3 GHz," New America Foundation, Apr. 2004.

[179] P. Marshall, "Unmet Challenges in the Technology to Exploit Dynamic Spectrum and Networking," in *First International Workshop on Technology and Policy for Accessing Spectrum Workshop at the IEEE Wireless Internet Conference*, Boston, MA, Aug. 5, 2006.

[180] P. Marshall, "Spectrum Management and Other Ongoing Cross-Layer DARPA Initiatives," in *National Science Foundation NETS Focus Area Symposium*, Arlington, VA, Feb. 5, 2004.

[181] P. Marshall, "Connectionless Networks: Radios and Networking Technology for Sensor Applications," in *Conference on Unattended Ground Sensor Technologies and Applications VII, International Society for Optical Engineering Symposium on Defense and Security*, Orlando, FL, Mar. 29, 2005.

[182] V. Cerf and R. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Transactions on Communications*, vol. COM-22, No. 5, pp. 637-648, May 1974.

[183] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the Self-Similar Nature of Ethernet Taffic (extended version)," in *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1-15, Feb. 1994.

[184] P. D. Sutton, K. E. Nolan, and L. E. Doyle, "Cyclostationary Signatures in Practical Cognitive Radio Applications," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 1, Jan. 2008.

[185] P. Sutton, P. Ozgul, K. E. Nolan, and L. E. Doyle, "Bandwidth-Adaptive Waveforms for Dynamic Spectrum Access Networks," in *3rd IEEE International Symposium on Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.

[186] S. Parikh and R. Durst, "Disruption Tolerant Networking for Marine Corps CONDOR," in *Proceedings of the 2005 Military Communications Conference*, Atlantic City, NJ, Oct.17-20, 2005.

[187] K. Scott, "Disruption Tolerant Networking Proxies for On-the-Move Tactical Networks," in *Proceedings of the 2005 Military Communications Conference*, Atlantic City, NJ, Oct. 17-20, 2005.

[188] L. Lehman, S. Garland, and D. Tenenhouse,"Active Reliable Multicast," in *Proceedings of IEEE INFOCOM '98*, San Francisco, CA, 1998.

[189] S. Symington, "Delay-Tolerant Networking Metadata Extension Block," Internet Draft draft-irtf-dtnrg-bundle-metadata-block-01, Feb. 17, 2009.

[190] L. Van der Perre, B. Bougard, J. Craninckx, W. Dehaene, W. Hollevoet, L. Jayapala, M. Marchal, P. Miranda, M. Raghavan, P. Schuster, T. Wambacq, P. Catthoor, and F. Vanbekbergen, "Architectures and Circuits for Software-Defined Radios: Scaling and Scalability for Low Cost and Low Energy," in *IEEE International Solid-State Circuits Conference*, San Francisco, CA, Feb. 11-15, 2007.

[191] M. McHenry, E. Livsics, A. Leu, D. McCloskey, V. Suri, D. Padrick, and A. Poddar, "Tuner Utilization and Feasibility (TUF) Study Final Report," Unpublished report to DARPA, available from Shared Spectrum Company or DARPA, Dec. 9, 2005.

[192] D. Cabric and R. Brodersen, "Physical Layer Design Issues Unique to Cognitive Radio Systems," in *IEEE 16th International Symposium on Personal, Indoor and Mobile Radio Communications*, 2005.

[193] B. Vigoda, "Analog Logic: Continuous-Time Analog Circuits for Statistical Signal Processing," Massachusetts Institute of Technology, Ph.D Dissertation Sept. 2003.

[194] M. Kokar, D. Hillman, S. Li, B. Fette, P. Marshall, M. Cummings, T. Martin, and J. Strassner, "Towards a Unified Policy Language for Future Communications Networks: A Process," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, 2008.

[195] P. Marshall, "Current R&D Investments in Dynamic Spectrum Technology," in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, 8-11 Nov. 2005.

[196] K. Feeney, D. Lewis, P. Argyroudis, K. Nolan, and D. O'Sullivan, "Grouping Abstraction and Authority Control in Policy-based Spectrum Management," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, IE, 2007.

[197] D. Lewis, K. Feeney, C. Tsarouchis, D. O'Sullivan, and J. Keeney, "Integrating the Policy Dialectic into Dynamic Spectrum Management Issues," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, IE, Apr. 22-25, 2007.

[198] F. Perich, R. Foster, P. Tenhula, and M. McHenry, "Experimental Field Test Results on Feasibility of Declarative Spectrum Management," in *3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, 2008*, Chicago, IL, Oct. 2008.

[199] Y. Guo, Z. Pan, and J. Heflin, "Lehigh University Benchmark (LUBM): A Benchmark for OWL Knowledge Base Systems," *Journal of Web Semantics*, vol. 3(2), pp. 158-182, 2005.

[200] M. McHenry. Personal correspondence with the author.

[201] C. Doerr, M. Colagrosso, D. Grunwald, and D. Sicker, "Scalability of Cognitive Radio Control Algorithms," in *3rd International Symposium on Wireless Pervasive Computing*, May 7-9, 2008.

[202] K. Takeuchi, T. Fukuhara, S. Nomura, and S. Yamamoto, "Cognitive Radio Using Multi-transmission Links: A novel approach effective in metropolitan areas," in *2nd International Conference on Cognitive Radio Oriented Wireless Networks and Communications*, pp. 291-297, Aug. 1-3, 2007,

[203] J. Redi, S. Kolek, K. Manning, C. Partridge, R. Rosales-Hain, R. Ramanathan, I. Castineyra, "JAVeLEN - An ultra-low Energy Ad hoc Wireless Network," *Ad Hoc Networks*, vol. 6, no. 1, pp. 108-126, Jan. 2008.

[204] Strategy Analytics, Inc., "Global Wireless Home Device Sales To Reach 314 Million Units By 2010," July 26, 2006, http://www. strategyanalytics.com.

[205] International Organization for Standardization, "ISO/IEC 13211: Information technology — Programming languages — Prolog," Geneva.

[206] A. J. Bonner and M. Kiffer, "Transaction Logic Programming," in *Proceedings of the 10th International Conference on Logic Programming*, Budapest, Hungary, June 1993.

[207] M. Kifer, G. Lausen, and J. Wu, "Logical Foundations of Object Oriented and Frame Based Languages," *Journal of the Association for Computing Machinery*, May 1995.

[208] I. Horrocks, P. F. Patel-Schneider, H. Boley, S. Tabet, B. Grosof, and M. Dean, "SWRL: A Semantic Web Rule Language Combining OWL and RuleML," World Wide Web Consortium member submission, http://www.w3.org/Submission/SWRL/.

[209] M. Cummings, S. Li, M. Kokar, B. Fette, B. Lyles, P. Marshall, and D. Hillman, "Activities of SDR Forum MLM Working Group on a Language for Advanced Communications Systems Applications," in *50th Software Defined Radio Forum General Meeting*, Boston, MA, 2008.

[210] L. Kent, K. Young, O. Younis, D. Shallcross, K. Sinkar, A. McAuley, K. Manousakis, K. Chang, and C. Graffand, "Network Science Based Approaches to Design and Analyze MANETS for Miltary Applications," *IEEE Communications Magazine*, vol. 46, no. 11, pp. 55-61, Nov. 2008.