A Day in the Life of the RF Spectrum

by

James E. Cooley

B.S., Computer Engineering
Northwestern University (1999)

Submitted to the Program of Media Arts and Sciences,
School of Architecture and Planning,
In partial fulfillment of the requirement for the degree of

MASTER OF SCIENCE IN MEDIA ARTS AND SCIENCES
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2005

Signature of Author: _____

Program in Media Arts and Sciences
August 5, 2005

Certified by: _____

Andrew B. Lippman
Senior Research Scientist of Media Arts and Sciences
Program in Media Arts and Sciences
Thesis Supervisor

Accepted by: _____

Andrew B. Lippman
Chairman, Departmental Committee on Graduate Studies
Program in Media Arts and Sciences

A Day in the Life of the RF Spectrum

by

James E. Cooley

Submitted to the Program of Media Arts and Sciences,
School of Architecture and Planning,
September 2005 in partial fulfillment of the
requirement for the degree of
Master of Science in Media Arts and Sciences

Abstract

There is a misguided perception that RF spectrum space is fully allocated and fully used though even a superficial study of actual spectrum usage by measuring local RF energy shows it largely empty of radiation. Traditional regulation uses a fence-off policy, in which competing uses are isolated by frequency and/or geography. We seek to modernize this strategy. Given advances in radio technology that can lead to fully cooperative broadcast, relay, and reception designs, we begin by studying the existing radio environment in a qualitative manner. We wish to objectively understand the purpose of a particular transmission, its threshold of allowable interference, and whether anyone is attempting to receive it. We wish to propose ways in which cognitive radio systems might coexist with legacy radio systems.

In Chapter 1, we review the conditions that led to the current regulatory climate. Chapter 2 discusses the purpose of this thesis and how the work done relates to cognitive radio technologies. Chapter 3 discuses the design of data capture and analysis modules used to better understand RF spectrum space usage. Chapter 4 applies the software modules to a range of spectrum space and evaluates the results.

Thesis Supervisor: Andrew B. Lippman
Title: Senior Research Scientist of Media Arts and Sciences

A Day in the Life of the RF Spectrum

by

James E. Cooley

Thesis Committee:

Advisor: _____

Dr. Andrew Lippman
Senior Research Scientist of Media Arts and Sciences
Massachusetts Institute of Technology

Reader: _____

Dr. David Reed
Adjunct Professor of Media Arts and Sciences
Massachusetts Institute of Technology

Reader: _____

Dr. Philip J. Fleming
Fellow of the Technical Staff
Motorola Inc.

Acknowledgements

The completion of this thesis would only be possible with the generous support from the following people:

My advisor, Dr. Andrew B. Lippman, whose boundless energy requires that his students innovate and learn at all times

Dr. David Reed who was my first contact at the Media Lab and whose musings on a wide-variety of topics inspires students… It was David Reed who sparked my interest in GNU Radio software-defined radio

Matt Ettus and Eric Blossom, the gurus of GNU Radio

My Viral Communications colleagues past and present: Marios, Jeff, Aggelos, Dean, Vyzo, Fulu, Hector, and Kwan

John DiFrancesco who quite rightly views Media Lab student construction projects with a healthy dose of skepticism!

Deb and Polly without whom, the details would get lost

To my parents, brothers, LouAnne, and niece and nephew

To Eileen

To all of my colleagues who made my time at the lab one the most rewarding learning experiences I've had

Table of Contents

List of Figures

Introduction

A Day in the Life of the RF Spectrum is undertaken as a study of radio frequency spectrum usage in the United States with the express hope that it will lead to the application of new radio technologies to improve radio frequency allocation schemes and maximize the benefits of various modern radio technologies.

Radio designs have evolved from huge, inefficient brute-force broadcasters to more-efficient digital systems to flexible software-defined radio systems. RF spectrum usage varies greatly with time, location, usage, purpose, and technologies employed. We begin in earnest by focusing our study on legacy analog systems such as television and FM Radio Broadcasts and their detection by computer systems. We ask how efficiently this space is used and inspire more-efficient use, re-use, or multi-use by describing ways to objectively qualify the nature of found signals. Rather than rely on an antiquated notion of fencing off radio systems from each other, we would like to inspire all new

radio systems to be built so that the nodes are inherently cooperative with each other and with legacy systems.

In order to accomplish this goal, software modules to classify and analyze radio signals were built for this thesis using an open-source software radio platform. These modules may be used as a basis for building a wholly-automatic cooperative radio system of the future: a system in which each node can understand its RF environment in a *qualitative* way, cooperate with its peers, and provide wireless services for a wide variety of applications. Most importantly, it will open RF spectrum space for public use not in such a way as to create a chaotic free-for-all, but in order to inspire a wireless communications revolution comparable to the Internet boom. While we don't cover the entire spectrum in this thesis, we hope to inspire a new way of thinking about the RF spectrum as a whole rather than as a sum of many pieces.

This thesis is divided into four major chapters. First, we examine the history of regulation in the United States and how legacy affects our policies today. Next, the concept of cognitive radio is discussed and how the work done for this thesis directly applies to this technology. The software-defined radio architectures designed and used for study are then detailed. The designs outlined here are intended to be a basis for building better-informed cognitive radio systems. Finally, we present the results of these studies, and the conclusion applies the results to the overall broad goals outlined here.

1 Background

This chapter discusses the history of radio technologies, how they have influenced regulation, and how modern radio systems inspire a new direction in thinking about regulation policy. This thesis addresses a way in which integration of modern systems with legacy technology may be possible.

1.1 History of Radio Designs

A study of modern RF spectrum usage in the United States is indelibly linked to history of radio system designs and regulation which followed. Spectrum allocation policy has traditionally followed a "fence-off" frequency allocation scheme, whereby channel usage is based on a clear separation in frequency space or geography. This is a transmitter-centric view of broadcast radio and can be seen as having evolved directly from old, inefficient, outdated radio designs which radiate at high power without sensing or adapting to neighboring transmitters or receivers. These antiquated systems were a

brute-force answer to a problem not nearly as well understood and engineered for as it is today.

1.1.1 Marconi and the Earliest Days of Radio

In the very earliest days of radio, in what must surely be one of the first live events covered by remote wireless technology, Guglielmo Marconi famously telegraphed reports from the America's Cup Yacht races of 1899 to the New York Herald offices. It was not long before he had competition. For the race of 1900, Lee de Forest challenged Marconi's monopoly on wireless. "To the great embarrassment of Marconi's engineers the two systems interfered with each other to such an extent that they had to agree with de Forest to transmit alternately at five-minute intervals."[1] This may be the first instance of cooperation among radio systems, but also one of the first instances of contention.

In the America's Cup case, neither Marconi nor de Forest had any notion of tuning or channelization because they were using spark-gap radios – literally powerful current arcs - which effectively radiated at all frequencies. The way in which the earliest pioneers achieved greater distances was by building larger spark-gaps and antennas. The spark-gap generator at Poldhu Station designed by Sir Ambrose Fleming for Marconi was particularly noteworthy and dangerous, as it was intended for transatlantic transmission. Arthur Blok worked for Fleming and described the spectacle, "...the roar of the discharge could be heard for miles along the coast...Every metal gutter, drainpipe or other object about the sheds on the site resonated freely and there was a minor chorus of ticks and flashes consonance with the discharge."[2] Figure 1.1 is a photograph of the station at Poldhu and Figure 1.2 shows some of the transmission equipment, including the infamous spark-gap generator and plenty of warning signs.

Today's digital cellular systems and other digital radios have given us so many options for multiple access that we now take the idea for granted. But, for most of broadcasting history, the only way to have multiple channels has been to separate them by frequency and/or in space. Separation by space also implies power considerations because stations that broadcast with more power generally cover more geographic area. Separation by frequency was soon discovered and will be discussed next. These considerations framed the evolution of spectrum policy.

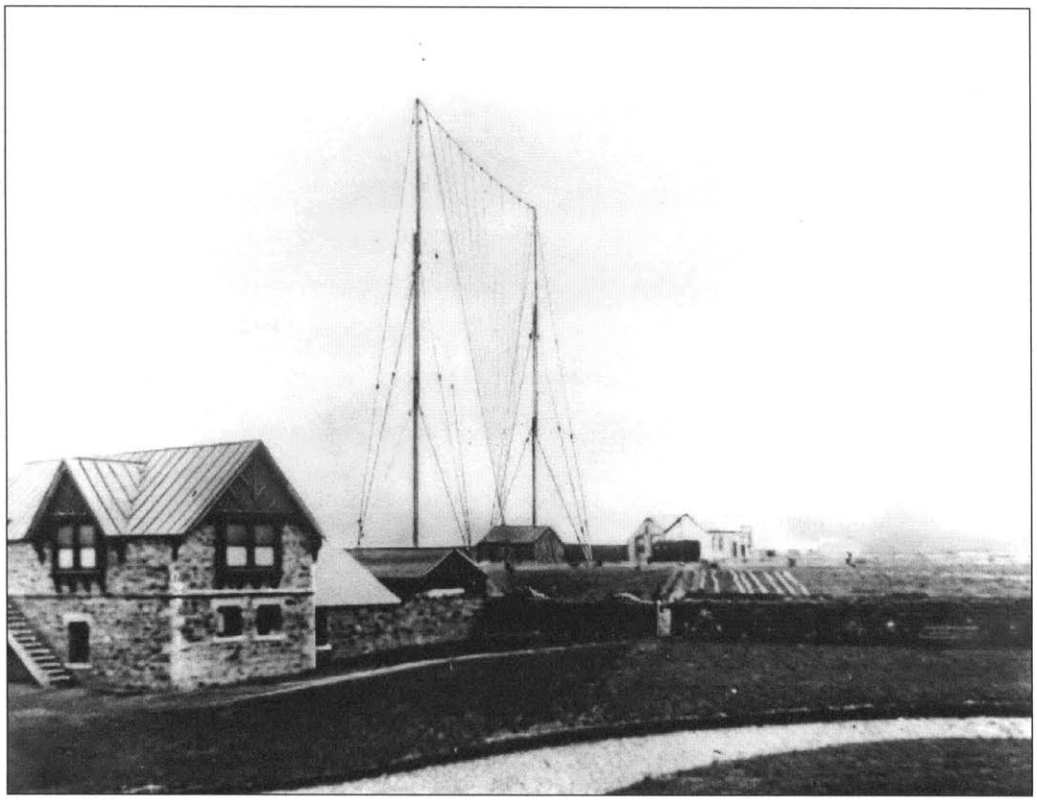

Figure 1.1 Marconi's Second Antenna at Poldhu, England, 1901[3]

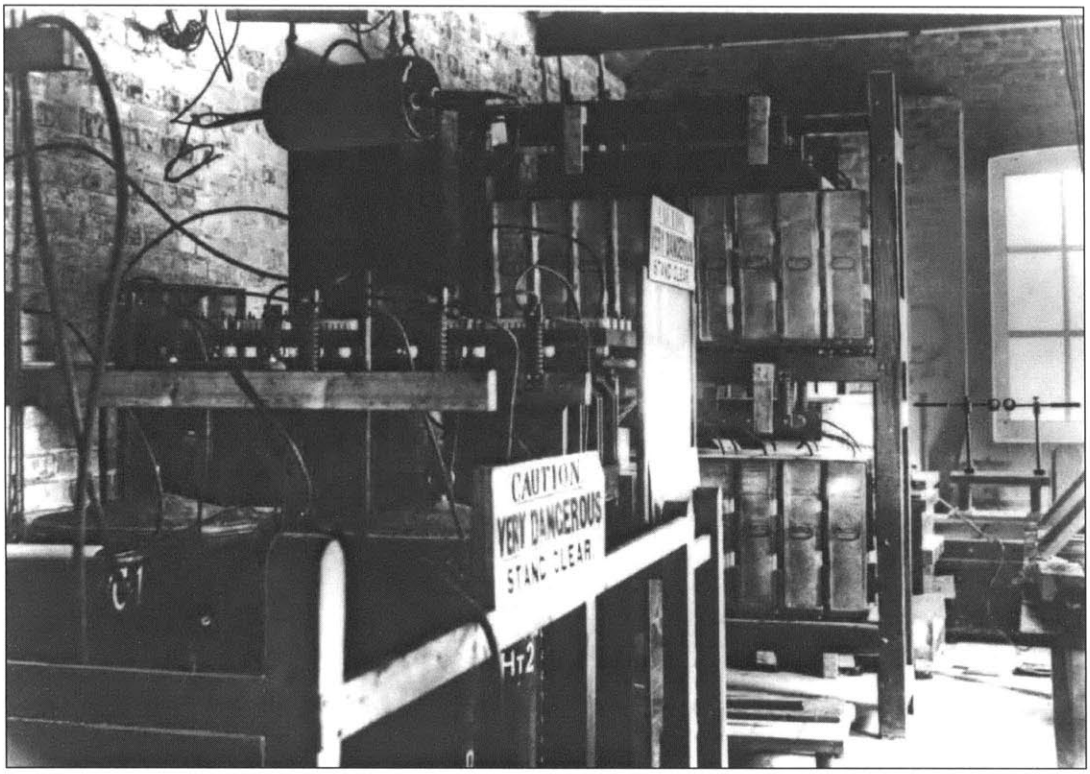

Figure 1.2 Spark-Gap Generator and Other Equipment at Poldhu Surrounded by Warning Signs[4]

1.1.2 Tuning, Channelization by Amplitude Modulation

If Marconi is the father of radio, Reginald Fessenden is surely the father of broadcasting. Moving beyond spark-gap transmissions on all frequencies, the invention of Amplitude Modulation allowed more than one station to transmit at a time by modulating each on a different frequency. The Alexanderson Alternator and Reginald Fessenden's heterodyne principal were key to this advance. Fessenden was chiefly interested in transmitting speech on a continuous wave as opposed to Morse Code by spark-gap. The powerful, 100,000-cycle alternators that Fessenden got from Alexanderson at General Electric served as the oscillators that would make this modulation possible.[5] Later these mechanical oscillators were replaced by the vacuum tube and still later, the transistor.

1.1.3 The Rise of Broadcasting, RCA, and Pre-Regulation

As radio technologies advanced, an understanding of what this new technology might be good for slowly grew out of the early point-to-point transmissions. Regulation in the United States evolved through a complicated interaction between technological innovation, social engineering, commercial interests, wartime powers, politics, and the changing role of government in radio activities.

The rise of broadcasting in the United States cannot be discussed without considering the role the Radio Corporation of America played in its development. During World War I, wireless was increasingly commandeered and used for ship-to-shore navy communications. Arrangements with American Marconi (owned 20% by its British parent company) worked out well during the war. Afterwards, the U.S. government and American business leaders worried that the British had a monopoly of this new technology on American soil, particularly given its military importance. RCA was chartered on October 17, 1919 when General Electric purchased the British shares of American Marconi, making it wholly American-owned.[6]

Licensing began as companies like RCA under the leadership of David Sarnoff became convinced that commercial broadcasting was far more lucrative than wireless telephony that had been the focus thus far. Manufacturers like Westinghouse saw broadcast radio as a means to sell receivers and they setup small stations accordingly.[7] As Sarnoff assumed the general manager position at RCA, he was able to influence electric firms and part-RCA owners Westinghouse and General Electric to produce smaller, inexpensive radio receiver sets. Growth of receiver sets (and therefore number of broadcast stations) was explosive. In January 1921 there were 36 stations nationwide. By

the end of 1922, there were 576.[8] New wireless broadcasters and stations, industrial equipment manufacturing giants like GE and Westinghouse, and network owners such as AT&T over whose wires many affiliate stations were "linked" formed a complicated marriage of the high-tech companies of the day.

Recognizing the importance of wireless and the ensuing chaos of many interests, then Secretary of Commerce Herbert Hoover convened in 1922 the first of four radio conferences to be held over the next few years. There were competing ideas and interests from diverse domains to resolve. For instance, who owned the radio spectrum? What role was government to play? How much threat did industry giants (such as the radio networks, RCA, AT&T and others) pose to smaller, independent broadcasters? What obligations did broadcasters have to provide public services? Did manufacturing and patent agreements between companies violate antitrust laws? How could patent disputes (of which there were a great many) be resolved? Was free speech guaranteed on the air waves? These are some of the questions brought up during the conferences and helped frame later discussions and define the future role of the FCC.[9] Much to the benefit of broadcasters like RCA who were already well entrenched, it was generally agreed that broadcasting would continue to be kept privately-owned and rf spectrum would belong to the public under government stewardship. From the fourth conference in 1925, delegates urged Congress to create a new regulatory body to manage all of this, similar to the Interstate Commerce Commission.

1.1.4 Regulation of Radio and Television

Regulation of broadcasting in the United States evolved through a complicated interaction between technological innovation, social engineering, commercial interests, wartime powers, politics, and the changing role of government in radio activities.

The Communications Act of 1934 established the FCC with real authority to license wireless communications as well as telephone and telegraph[10]. At the time, and the way that the technology operated, it is understandable from a technical standpoint that the only way to isolate powerful transmitters from each other was either to isolate geographically (the strength of a transmitted signal declines over a radiated distance) or to isolate by assigning non-interfering frequencies. The FCC's regulatory decisions in the intervening years had more to do with standardizing a particular technology over another than what frequency bands were used.

The relationship between the FCC and broadcasters ebbed and flowed over the years, having much to do with the monopoly position RCA was building. RCA had started the National Broadcasting Company in 1926 which had now grown huge with two major networks known as the "Red" and the "Blue." By 1927, NBC affiliates made up 25 percent of all broadcast stations and this number kept increasing.[11] The NBC Blue network was sold in 1943 and became the American Broadcasting Company after RCA faced antitrust pressures.

New technologies like FM and Television soon began to develop. Frequency Modulated broadcasts began in the 1930s. Not only did the technology deliver better audio fidelity than AM, it gave the industry what the amounted to a "second chance" in the minds of the FCC to offer a wide variety of programming such as news, music, and

dramas and an opportunity to once-again fuel revenue by radio set sales, not from the commercialism that had come to drive AM broadcasts. Initially, FM growth was steady, and the FCC assigned what would have been television channel 1 to FM.[12] NBC's WNBT became the first commercial station in 1940, and the FCC granted a few experimental licenses, but real progress was stalled until after World War II.

Industry battles ensued as the dominant companies vied for their version of television to be officially adopted and endorsed by the FCC. RCA had a monochrome system operating on VHF frequencies while CBS had what some consider a technologically superior color system operating on UHF. Although the FCC appeared to initially lean towards the RCA design, they recognized the superiority of the CBS version. They opened hearings on color systems in 1949 and CBS was allowed to begin color television broadcasts. The National Television Systems Committee (NTSC) was established as an industry group made up of RCA, GE, Sylvania, and others to directly lobby against this CBS design.[13]

Though it all, the RCA system eventually triumphed. RCA-owned NBC brought their biggest talents from radio to television and improved their own NTSC color system. The CBS technology was in trouble as receivers sales were down and color broadcasts were halted during the Korean War. While the 1950 decision in favor of NTSC was a coop for RCA despite the fact that it took a long time for lagging sales of the more-expensive receivers to increase, the decision to force receiver manufactures to include UHF and VHF frequencies opened the door to broadcast competition in 1952.[14]

Since then, FM and NTSC Television broadcasts have changed little with regard to regulation, although the next major broadcast changes are expected to come from

HDTV and DBS. This thesis is primarily concerned with how cognitive radio systems might adapt to these existing technologies and make more efficient use of the rf spectrum space allocated to them.

1.2 The Current Regulatory Environment

The Communications Act of 1934 established the Federal Communications Commission in part with a mandate to allocate radio frequency spectrum in the United States, and we have seen how this role has been played out with AM, FM and Television broadcasting.

Sharing spectrum management duties with the Office of Spectrum Management within the U.S. Department of Commerce, the two agencies have designated RF spectrum for various uses. Figure 1.3 shows the allocation chart as published by the Office of Spectrum Management[15]. This shows a veritable jungle of competing technologies, interests, restrictions and so forth and is only legible when printed as a poster! Each colored block represents a radio space allocated for a different use. Figure 1.4 is a small, readable block of the chart showing allocations for radio astronomy, aeronautical navigation, television channels, and FM Radio broadcasts.

Figure 1.3 The Published Frequency Allocation Chart

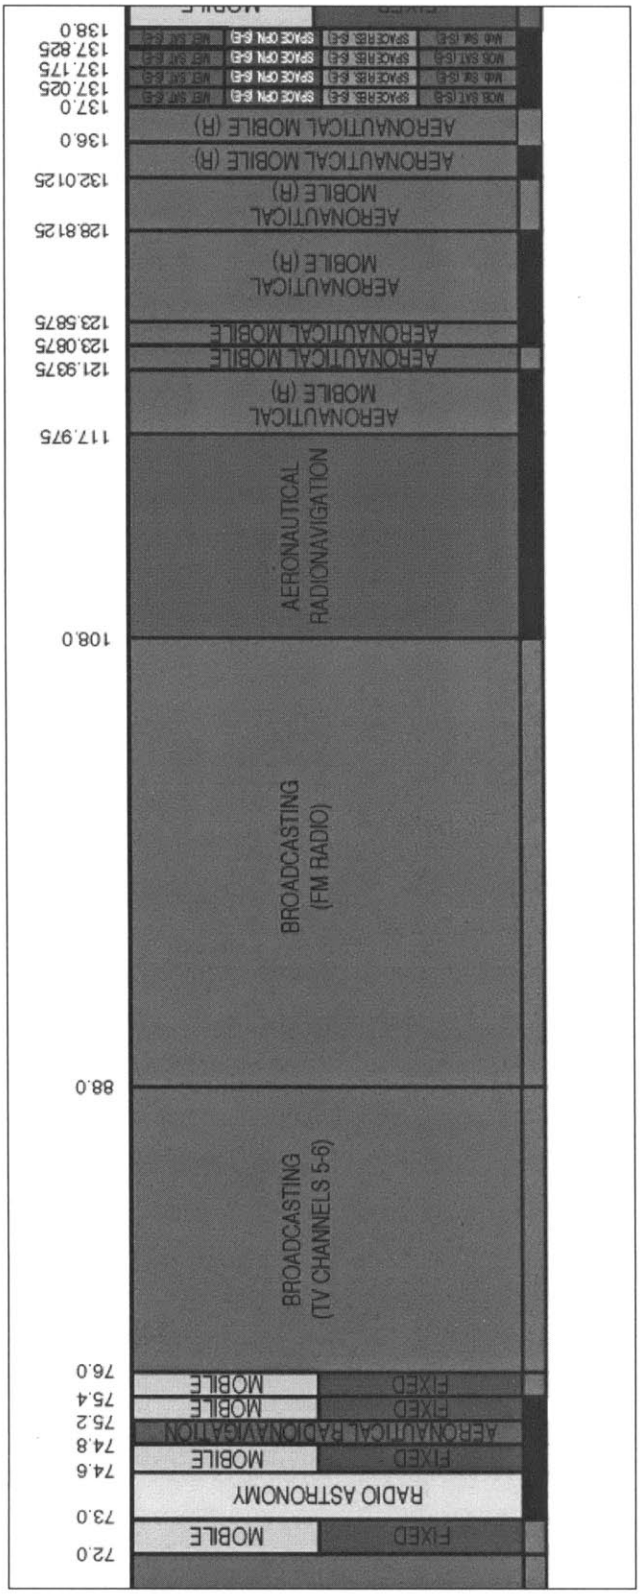

Figure 1.4 A Subsection of the Frequency Allocation Chart

These charts suggest a severe problem of scarcity. In this thesis, frequencies from 55MHz to 860MHz were studied. This range includes frequencies used for broadcast television, FM radio, aeronautical navigation, public safety (such as police and fire bands), and weather radio. For instance, radio frequencies from 54 MHz to 72 MHz are reserved for television broadcasting stations 2 through 4. The rest of the allocation map is similarly "designated" for all manner of uses. Some ranges are mapped to multiple uses, suggesting a desperate need to further partition a dwindling natural resource. In general, a naïve glance at RF frequency allocation from 3 KHz to 300 GHz implies virtually all RF spectrum space has now been allocated.

With scarcity comes trouble, as the disastrous government auctions of third-generation cellular phone space in Europe is a prime example. Prices for 3G spectrum space in the UK and Germany famously skyrocketed. Various reasons for the skyrocketing prices are suggested, but Cambridge economist Paul Klemperer disputes the suggestion that the cellular providers in the UK were unintentionally swept into the game. "[I]t is hard to imagine any company voluntarily paid billions of dollars for a 3G license. But volunteer they all did. Not only that but they celebrated their victories. The stock market was happy, and shortly after the UK auction Huthinson even resold part of its license at a profit."[16] There are many causes of these disasters and ways to understand the problem from multiple viewpoints. Clearly, owning the exclusive right to this dwindling frequency space was seen as a very valuable property.

The underlying problem from an engineering standpoint is the perception of scarcity. There should not need to be large auctions because the "scarcity" is caused by regulation based on politics and dated technical designs.

Regulatory agencies take a transmitter-centric view. Given modern radio technologies (software-defined radio, cognitive radio, ad-hoc wireless schemes) a receiver-centric scheme could be designed. More important than whether a transmitter is present and broadcasting in a given locality is whether there are any receivers trying to receive a transmitted signal. For instance, in-building reception of over-the-air television broadcasts is often abysmal. (Certainly, it is within the Media Lab). So, one can fairly reasonably guess that nobody is trying to watch television within the building because they *cannot.* There is no reason to prevent reuse of these frequencies for other uses. A highly periodic signal, such as one used to open a garage door, is idle most of the time. A scheme could be setup to reuse the frequencies when not needed to open garage doors. A more controversial example would be to reuse frequencies used in emergency communications, which are also idle more often than not. A system engineered to reuse these frequencies would need to guarantee emergency access as necessary.

1.3 Living With the "Full Spectrum" and Legacy Spectrum Usage

Despite a perception that RF spectrum space is fully allocated and fully used, recent studies of actual spectrum usage by measuring local RF energy clearly show it largely empty of radiation.

We find ourselves in an age in which unlicensed spectrum space is dwindling. In the FCC Workshop on Cognitive Radios in 2003, the Shared Spectrum Company presented a number of conclusions about uneven spectrum usage in practice. In general, they found that "Many bands have no detectable occupancy," "Some bands have low occupancy," and "Some bands have high occupancy."[17]

The FCC is keenly interested in new technologies which might mitigate this problem. DARPA's XG research program seeks to build a "Set of Advanced Technologies for Dynamic Spectrum Access" in part by taking "Temporal, Spectral, Dimensional" and energy measurements.[18]

1.3.1 Cognitive Radio Systems and Radio Resource Management

First defined in his doctoral dissertation, Joseph Mitola defines "cognitive radio" broadly as combining model-based reasoning with software-defined radio. In such a system, a user's "communications needs" are judged according to use context and "radio services and wireless services most appropriate to those needs" are provided.[19]

Of the use-cases cognitive radio systems are said to cover, "Spectrum Pooling" most closely relates to the goals for efficient use investigated in this thesis. Mitola outlines a radio "etiquette" whereby occupiers/owners of spectrum space negotiate a brief rental period of that space with a potential user. Mitola describes this as "... the set of RF bands, air interfaces, protocols, spatial and temporal patterns, and high level rules of interaction that moderate the use of the radio spectrum. Etiquette for spectrum pooling includes the spectrum renting process, assured backoff to authorized legacy radios, assured conformance to precedence criteria, and order-wire network, and related topics"[20]

1.3.2 The FCC and Cognitive Radios

The FCC has accepted the idea of cognitive radio technology being used to mitigate the scarcity problem. In March, 2005, Docket No. 03-108, the FCC issued regulatory rule changes aimed at cognitive radios. Figure 1.5 lists areas in which the FCC sees cognitive radio as helping ease the spectrum scarcity problem.[21]

- ***Frequency Agility*** - the ability of a radio to change its operating frequency to optimize use under certain conditions
- ***Dynamic Frequency Selection (DFS)*** – the ability to sense signals from other nearby transmitters in an effort to choose an optimum operating environment
- ***Adaptive Modulation*** – the ability to modify transmission characteristics and waveforms to exploit opportunities to use spectrum
- ***Transmit Power Control (TPC)*** – to permit transmission at full power limits when necessary, but constrain the transmitter power to a lower level to allow greater sharing of spectrum when higher power operation is not necessary.
- ***Location Awareness*** - the ability for a device to determine its location and the location of other transmitters, and first determine whether it is permissible to transmit at all, then to select the appropriate operating parameters such as the power and frequency allowed at its location.
- ***Negotiated Use*** - a cognitive radio could incorporate a mechanism that would enable sharing of spectrum under the terms of a prearranged agreement between a licensee and a third party. Cognitive radios may eventually enable parties to negotiate for spectrum use on an ad hoc or real-time basis, without the need for prior agreements between all parties.

Figure 1.5 Features of Cognitive Radio Systems Seen As Mitigating Spectrum Disuse by the FCC

In 2004, the FCC proposed rules for allowing use of "unused" television spectrum frequencies. Among the techniques required to be used by devices outlined in Dockets 04-186 and 02-380 concern are "smart radio features" that "employ spectrum sensing

techniques that would determine if the signals of authorized TV stations are present in the area."[22]

1.3.3 The FCC Noise Temperature

Central to the FCC's understanding on how advanced cognitive radio systems might be able to help the spectrum scarcity problem is the concept of noise temperature. Figure 1.6 Illustrates the FCC Concept of Noise Temperature. They refer to the existing licensee as the "primary licensee" and any others as "secondary." The principle means of multi-use would be to make sure that secondary users operate within a range that the primary licensee sees as noise and therefore inconsequential.

In this thesis, we propose that cognitive radios should build and use models more complete than those built on interference temperature alone.

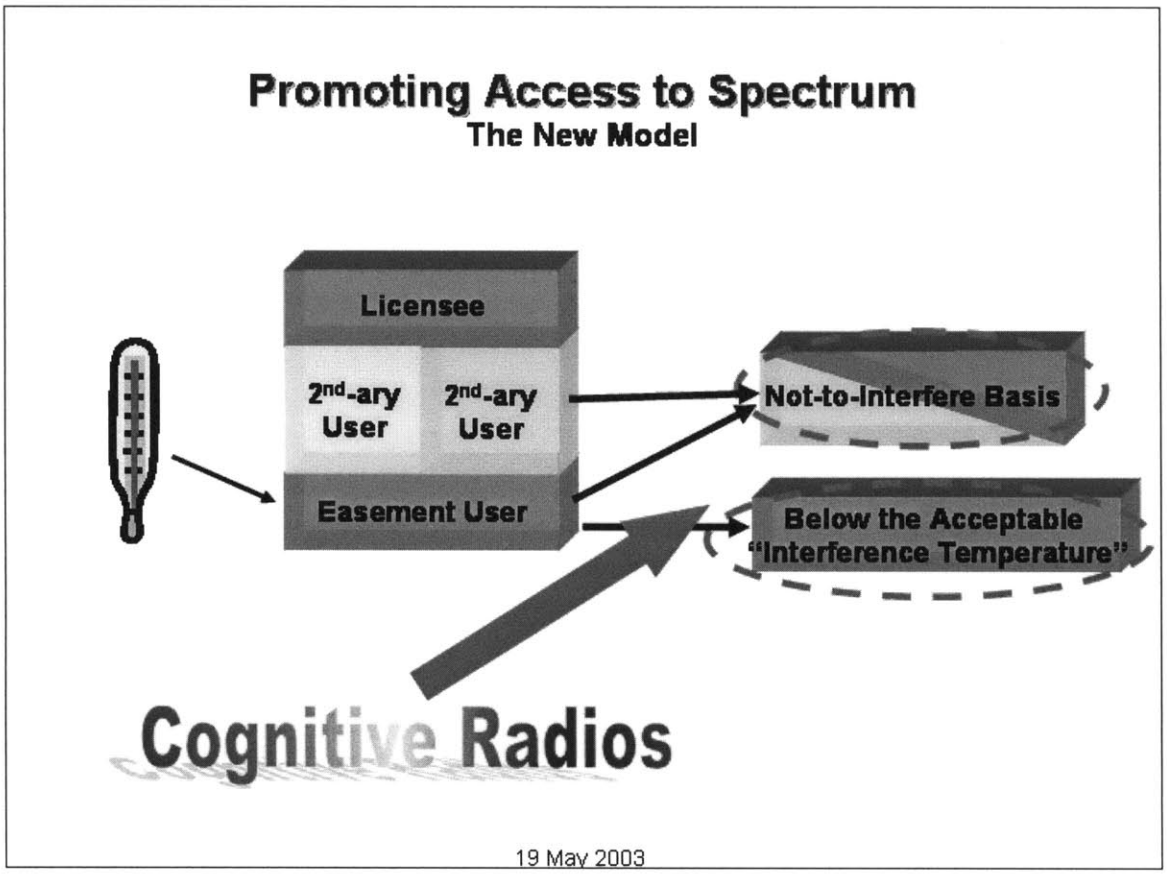

Figure 1.6 The FCC Concept of Noise Temperature[23]

2 A Day in the Life of the RF Spectrum

This thesis seeks to introduce application of cognitive radio techniques to understand legacy usage beyond energy measurements. We want to build an informed model of spectrum usage that can be used to create better cognitive radio systems.

Here, we propose the idea of a cognitive radio system sensing its radio environment, building a model of existing occupants, and using that model to more-efficiently use the radio frequency spectrum.

In his dissertation, Mitola outlines several use-case domains in which cognitive radio systems operate. As shown in Figure 1.1, the work done for this thesis falls under the "Radio Resource Management" use-case outlined by Mitola.[24]

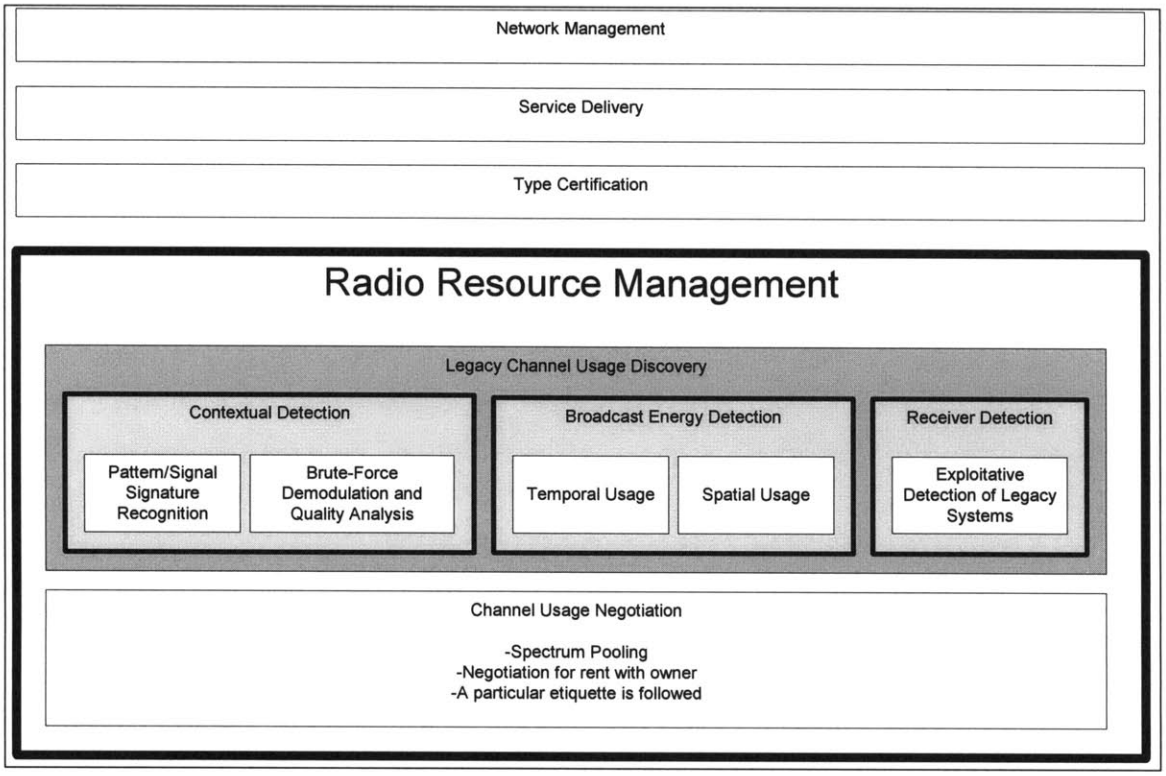

Figure 2.1 Cognitive Radio Use-Cases and A Day in the Life of the RF Spectrum

As discussed earlier, Mitola does address the issue of spectrum scarcity mitigation by proposing a space-leasing protocol or "etiquette" which might be used.

2.1 Legacy Channel Usage Discovery

This thesis is concerned with a more fundamental task within the Radio Resource use-case: Legacy Channel Usage Discovery. To better build a model of the surrounding RF environment, it is proposed that cognitive radio systems need a way to map and understand the existence of legacy or otherwise unknown radio systems. It is important to be armed with this information because a better judgment on how to navigate legacy space can be made. We break down this need into three sub-categories.

2.1.1 Contextual Detection

Contextual Detection is concerned with understanding the context of an unknown signal. For instance, if a particular modulation scheme can be identified, methods to best avoid disruption of that scheme can be employed. NTSC television modulation is an example of a modulation scheme in which much time is wasted between subsequent scan frames. If a cognitive radio system could understand that an unknown signal is NTSC, it could apply a modulation scheme appropriate to exploit the wasted time. We demonstrate two general types of detection which might be performed.

First, given an unknown signal, we can perform signal pattern recognition techniques in order to attempt to identify its nature. NTSC is a classic example of a modulation scheme which has a very identifiable signature in the frequency domain and we use this fact to develop algorithms to detect it.

Second, the flexibility of software radio platforms to "become" whatever radio system is desired through a simple software switch inspires a brute-force method of detection. In this scheme, the unknown signal is demodulated and the quality of the output is evaluated in order to infer whether the proper demodulation scheme was used.

2.1.2 Broadcast Energy Detection

In the more traditional and limited studies of radio spectrum usage so far, simple energy measurements have been done. Armed with even this type of knowledge, a cognitive radio system can make a more informed model of its environment. In this thesis, basic energy measurements were used to study temporal and spatial usage.

2.1.3 Receiver Detection

Beyond the idea of detecting broadcasts in rf spectrum space, receiver detection is an attempt to understand if there are any *listeners* of a given signal. If a cognitive radio system could understand that there is nobody currently within its coverage area *listening* to a particular broadcast, the system might be able to reuse that space. Figure 2.2 illustrates this idea. Although the two cognitive radios are within the coverage area of a broadcast system, there is no receiver in their immediate vicinity (indicated by the dotted regions).

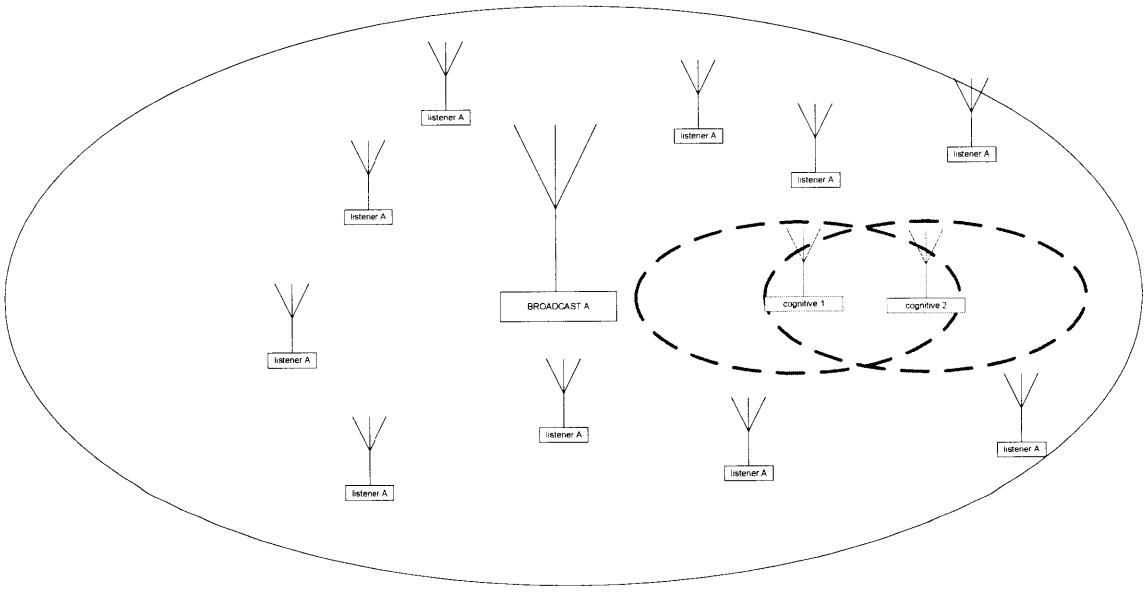

Figure 2.2 Cognitive Radios Detecting No Listeners in their Coverage Area

Future receiver designs would ideally have some announcement capability whereby they could transmit a small amount of information including what they are listening to. In this thesis, we exploit aspects of the design of legacy receivers in order to try to detect them.

3 Method

This chapter describes the motivations behind and designs of the software modules that were built for this thesis as well as background on the GNU Radio software and Universal Software Radio Peripheral (USRP) platforms. First, the evolution of the software architecture used for this thesis is described. Next, background information on the GNU Radio and USRP platforms is detailed. Finally, core C++ GNU Radio modules and Python building blocks designed specifically for this thesis are presented.

3.1 Architecture Evolution

The original hope for this thesis was to build a system capable of doing capture and real-time analysis. The idea was to be able to direct the software to examine a given frequency space and report on the content found at that space and correlate the findings with time and physical location. The size and complexity of such an architecture is

beyond what the USRP, GNU Radio software and basic personal computers can today support, so instead I have focused on developing techniques and algorithms to identify existing signals.

The first architecture designed on this principal was an all-in-one capture design in which raw samples were captured and correlated with location data and analyzed at a later time. The GNU Radio platform is able to receive raw sample data over the USB interface and then save it to a binary file which can later be examined and processed as if it were real-time data. In support of this architecture, specialized GNU Radio modules were built to save GPS location and time metadata along with raw sample data and allow safe multithread read/write access. Although not actually used in the final architecture designs used to produce data for this thesis, these software blocks are described in Section 3.3, "Spectrum Analysis Core Library Modules" below and nonetheless demonstrate a novel usage of software radio in that they correlate raw radio energy with location data.

Given the limitations of the USB interface, space requirements and processing power, it proved too difficult to reliably capture enough data at once to make this a viable solution for capturing and studying RF spectrum usage. For instance, an NTSC television station is 6MHz wide. This system is not able to record a bandwidth of 6MHz in real time. I investigated compressing the data on the fly using gzip, but the processing needed for this compression severely compromised the system's ability to receive and process incoming data from the USRP and write it to disk. A multiprocessor system might be able to perform better in this regard, but for the goals of this thesis, the architecture design shifted from a broad, all-encompassing design to more-practical software suites.

The final architecture presented here is a more moderate approach. In the software suites designed, an attempt was made to do as much processing in real-time as possible and reduce the amount of data saved. The capture applications were designed with the intended analysis of each application in mind. Though this does not fulfill the desire to build a comprehensive, real-time system that is ready to be used in a cognitive radio application, it does present the possibilities and methods of analysis which could be later employed.

The analysis modules presented here address legacy analog radio systems such as FM radio and NTSC television. From the fast-growing field of digital transmission schemes are technologies that were not studied here, but are no less important to the overall understanding of spectrum usage. Analysis of these digital signals is deferred to future work. When considering ways in which RF spectrum space may be allocated or used more efficiently, it seems best to first address older, less-efficient legacy systems. Despite the hype of ATSC (HDTV) broadcasts as replacing existing NTSC TV, it seems more likely that existing systems will remain operational at least for the foreseeable future. As the FCC and others have suggested, it therefore makes practical sense to build cognitive radio systems that can coexist with these systems to re-claim some of the most inefficiently-used space.

More advanced architectures are possible than those presented here. Using multiprocessing with several networked machines and concurrent analysis using several techniques is one scheme which may improve the processing power of a usage discovery system. Furthermore, more-specialized firmware and hardware may also be a way to address some of these issues in future work.

It should be clear from the following sections the power and flexibility software radio architectures can have. The fact that we are able to discuss radio system building blocks in terms of software modules that can be "connected" together using software represents a great technological advance over antiquated radio designs.

3.2 GNU Radio Software and Hardware Background

This thesis builds upon the GNU Radio software radio platform and the Universal Software Radio Peripheral (USRP) hardware. GNU Radio[25] is an open-source software-defined-radio project and the USRP[26] is a generic hardware device designed to operate with it.

The core of GNU Radio is written in C++ and provides a basic environment for building a pipeline combining software-based digital signal processing blocks. The "Simplified Wrapper and Interface Generator" (SWIG) is used to provide a simplistic Python interface and APIs for gluing dsp blocks together. Figure 3.1 diagrams this basic system.

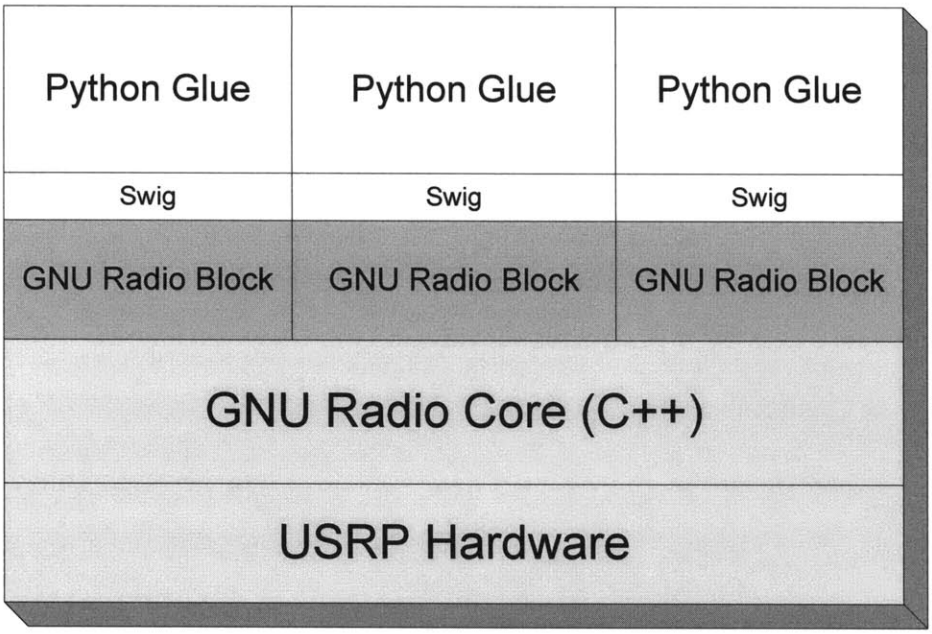

Figure 3.1 GNU Radio/USRP Core Block Diagram

There are three fundamental types of GNU Radio blocks: signal sources, filters, and signal sinks. For this thesis, the USRP was usually the signal source although sometimes a file source (a file with previously-recorded raw sample data in it) was used. There is a wide variety of filter blocks which perform operations on the input stream and produce the results on the output stream. Adders, multipliers, low-pass, conversion to/from number systems are all present as well as many other building blocks. Signal sinks can be files, or in most cases for the blocks implemented here, GUIs and Fast-Fourier Transforms or other special analysis blocks.

The USRP was outfitted with two "TVRX" daughter cards. These modules each consist of a Microtune 4937 cable tuner module which converts an input signal anywhere in the range 55 MHz – 850 MHz down to a center frequency of 5.75 MHz. Each output is then fed into a 64 Megabit A/D converter. From there, the data moves on to processing in the firmware of the on-board Altera Cyclone FPGA. The firmware is responsible for signal down-conversion to baseband frequency and decimation (keeping one in n sample,

which has the effect of reducing the overall sample rate). The data from the two input paths is multiplexed and output over the USB 2.0 port to be handled on the PC side. The USRP does have a transmit path too, but that was not used for this thesis.

It bears mentioning that the PC used for experiments in the thesis is a Dell Latitude D600 Laptop with a 1.6 GHz Intel Pentium M processor, 1 GB of memory, and running Debian (GNU) Linux using kernel version 2.6.10.

3.3 Spectrum Analysis Core Library Modules

Although not used in the final Python analysis suites discussed in the next section or for the results presented in the next chapter, a number of core GNU Radio library modules were built in support of this thesis and used in experiments leading to the final designs.

3.3.1 File Header Sink and Source

The gr_file_hdr_sink and gr_file_hdr_source building blocks are direct descendants of the GNU Radio core library file sink and source blocks, except that they offer the Python interface/glue application the chance to save metadata correlated along with each sample block. The read/write interfaces guarantee exclusive access to the metadata block by using posix semaphores to guard the header data structure. While the standard file sink and source blocks simply save chunks of raw sample data, a special file format was devised to accommodate this design. To simplify the design, each header item is stored as floating point numbers. The overall header size is listed first, followed by an arbitrarily-sized list of header items (which could be size zero). The number of header

items as well as the items themselves can be changed by the higher-level Python thread at any time in a thread-safe manner. Following the last header item (or header size if there are no header items), is stored the data block size. Following the data block size is a block of raw sample data as it might be stored in the standard file sink and source blocks. This format repeats once the block of raw sample data has been written. (The size of the raw sample data block is chosen internally in the GNU Radio pipeline). Figure 3.2 illustrates this format.

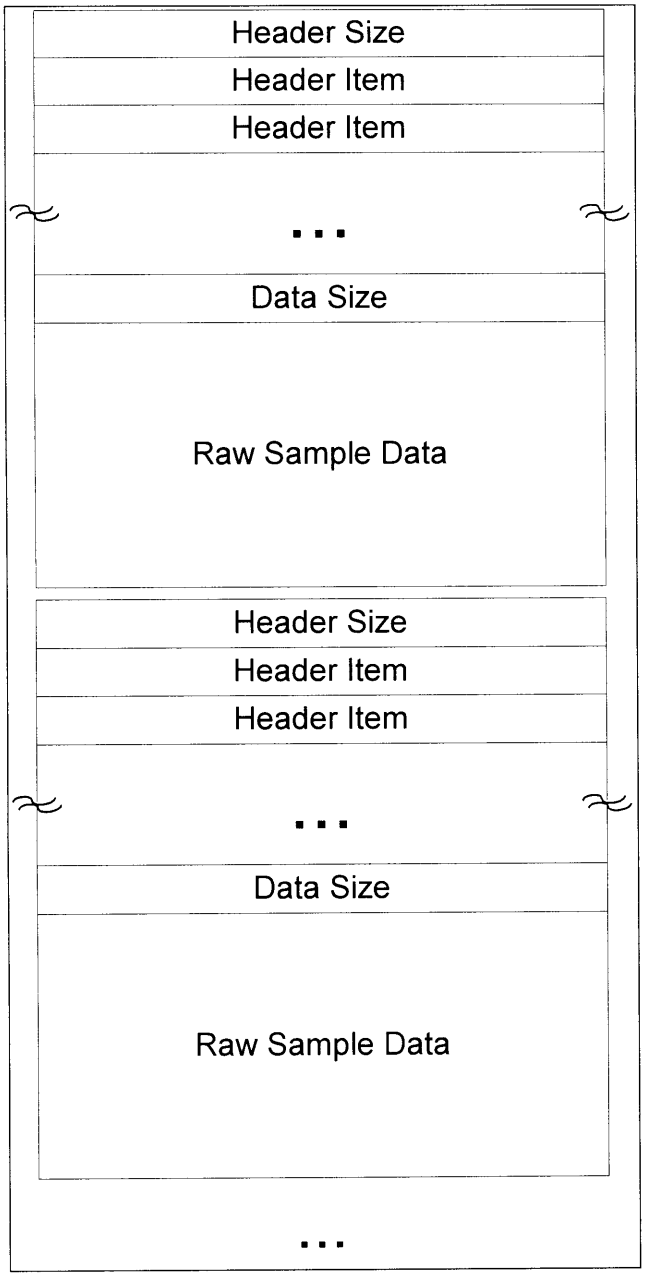

Figure 3.2 Data Format for Adding a Metadata Header to Raw Sample Data

These sink and source blocks were expressly intended for saving GPS location data correlated with raw sample data. Using the sink block, a Python thread queries the GPS unit and updates the header as necessary when the location changes. Later, when the

data file is read back, a Python thread monitors the location metadata for changes and the location associated with the sample data is updated accordingly.

3.3.2 FFT Reorder Block

GNU Radio uses the FFTW[27] library at its core in order to perform Fast Fourier transforms on sample data. This library outputs a data block of FFT buckets in which the 0^{th} and positive frequencies are first followed by the negative frequencies. The gr_cfft_reorder block alters this order such that the frequencies are output in order of increasing frequency. This block simplifies some analysis routines.

3.3.3 Sample Threshold Detector

This block is nearly identical to the simple_squelch block found in the GNU Radio core. A threshold level is set in the block and when an incoming signal sample magnitude equals or exceeds the level, a flag is set. This flag can be queried in a thread-safe manner from Python application code. In addition, the module outputs a binary stream of data in place of the input sample data: 1 for when a particular sample met the threshold and 0 for when it did not.

The original design of the Narrow Band FM usage study code used this module, but it later became more practical to build a simpler design in higher-level Python code.

3.4 Spectrum Analysis Python Suites

These high-level Python modules were designed specifically to capture meaningful data in support of goals of this thesis. The results of each capture are presented in the next chapter unless otherwise noted. There is some reuse of design blocks across the various modules. Targeting the design of each module towards a specific data capture need proved more practical than trying to design an all-in-one capture methodology.

3.4.1 Simple Capture

The Simple Capture Python suite grew out of the frustration with trying to create a complicated sample capture scheme in which raw sample data would be captured and later analyzed. In order to minimize the amount of data recorded at a given time, a Fast Fourier Transform of the data is taken in order to discover frequencies present during capture. The software is designed to gradually increment the tuned frequency in order to cover the range of the hardware, approximately 50MHz – 850MHz. The minimum, maximum, and average amplitude is recorded for each frequency bucket from the FFT over a few seconds time.

Two important workarounds were needed in this design in order to avoid limitations of the USRP hardware and firmware. First, the USRP uses a CIC filter in order to down-convert a desired frequency to baseband. The USRP has a sampling rate of 64MS/s and the data is decimated on the USRP by 16, thus giving an effective output rate

of 4MS/s. This introduces a curve in the baseline of the raw frequency samples as significant as 20dB from edge to center, as seen in the FFT graph of Figure 3.3.

In order to work around this problem, the solution is to further decimate the incoming data stream by 4 in the PC in order to minimize the effect of the curve. Since the decimated USRP output is 4MS/s this leaves a working data rate of 1MS/s. Figure 3.4 shows a flatter baseline after this is applied. This amounts to taking a one-quarter chunk of bandwidth from the center of the first graph. In the new smaller piece, the overall effect of the curve is far less extreme.

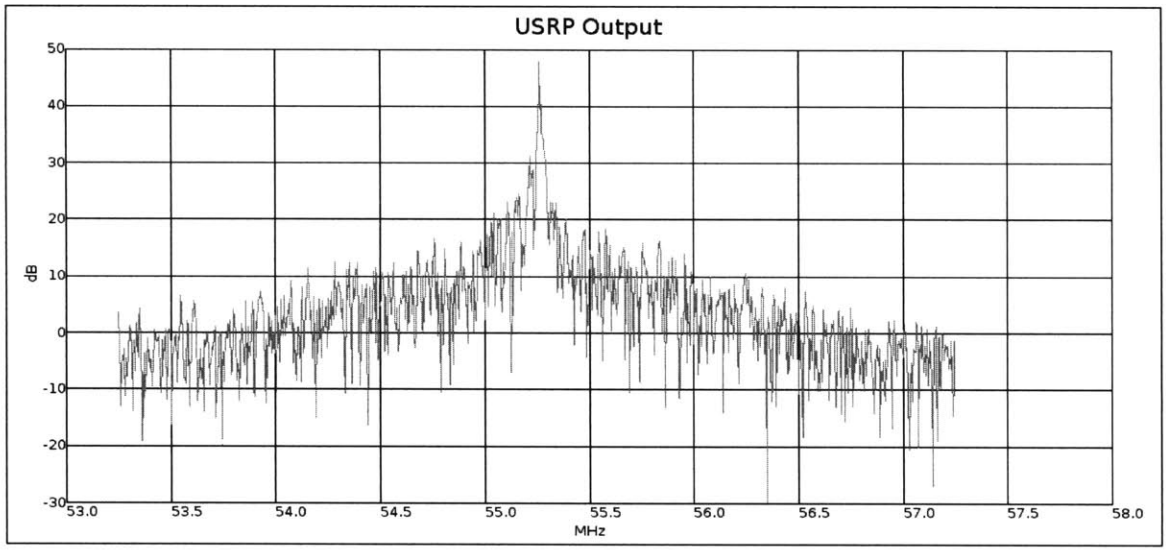

Figure 3.3 FFT With Curved Baseline As Output by the USRP CIC Filter

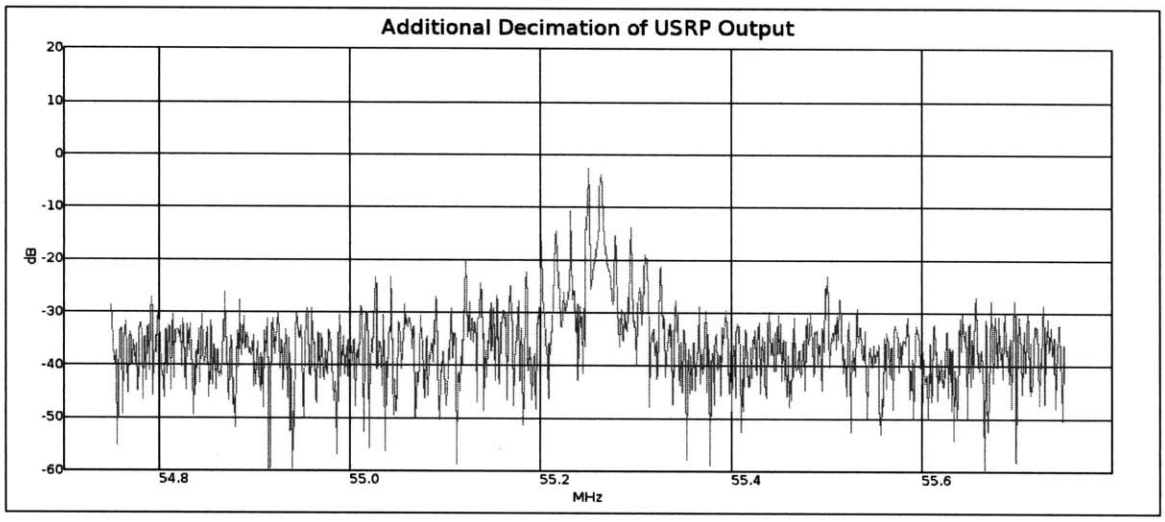

Figure 3.4 FFT With Flattened Baseline After Additional Decimation

The second workaround needed for the USRP is due to the fact that the hardware introduces a DC component at baseband 0 which is present in its output and appears as a spike in the FFT display. This spike is clearly shown in Figure 3.5, located on the graph at approximately 55.65 MHz.

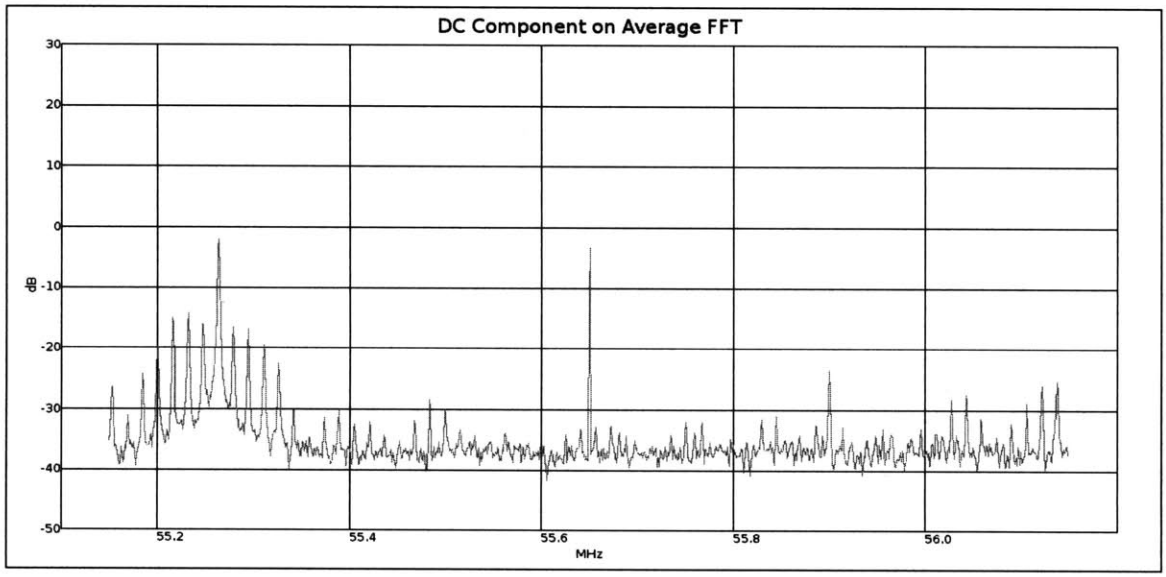

Figure 3.5 USRP DC Component Spike

Any analysis of spectrum data must account for this spike. In this simple capture suite, it was avoided altogether. Care was taken in how the USRP is tuned such that the DC spike always falls at baseband. Once the FFT is taken on the data, it is divided into four, evenly-spaced groups of 256 buckets each. The frequency data from the first and fourth bucket is recorded, while the data from the second and third (containing the DC component spike) is discarded.

After a set delay, the capture tuning is incremented by half of the total frequency bandwidth captured by the FFT, thus making sure that an average FFT value is taken and recorded for each frequency. The accumulated average FFT data is re-assembled into a continuous spectrum. Figure 3.6 shows a diagram of this data interleaving scheme.

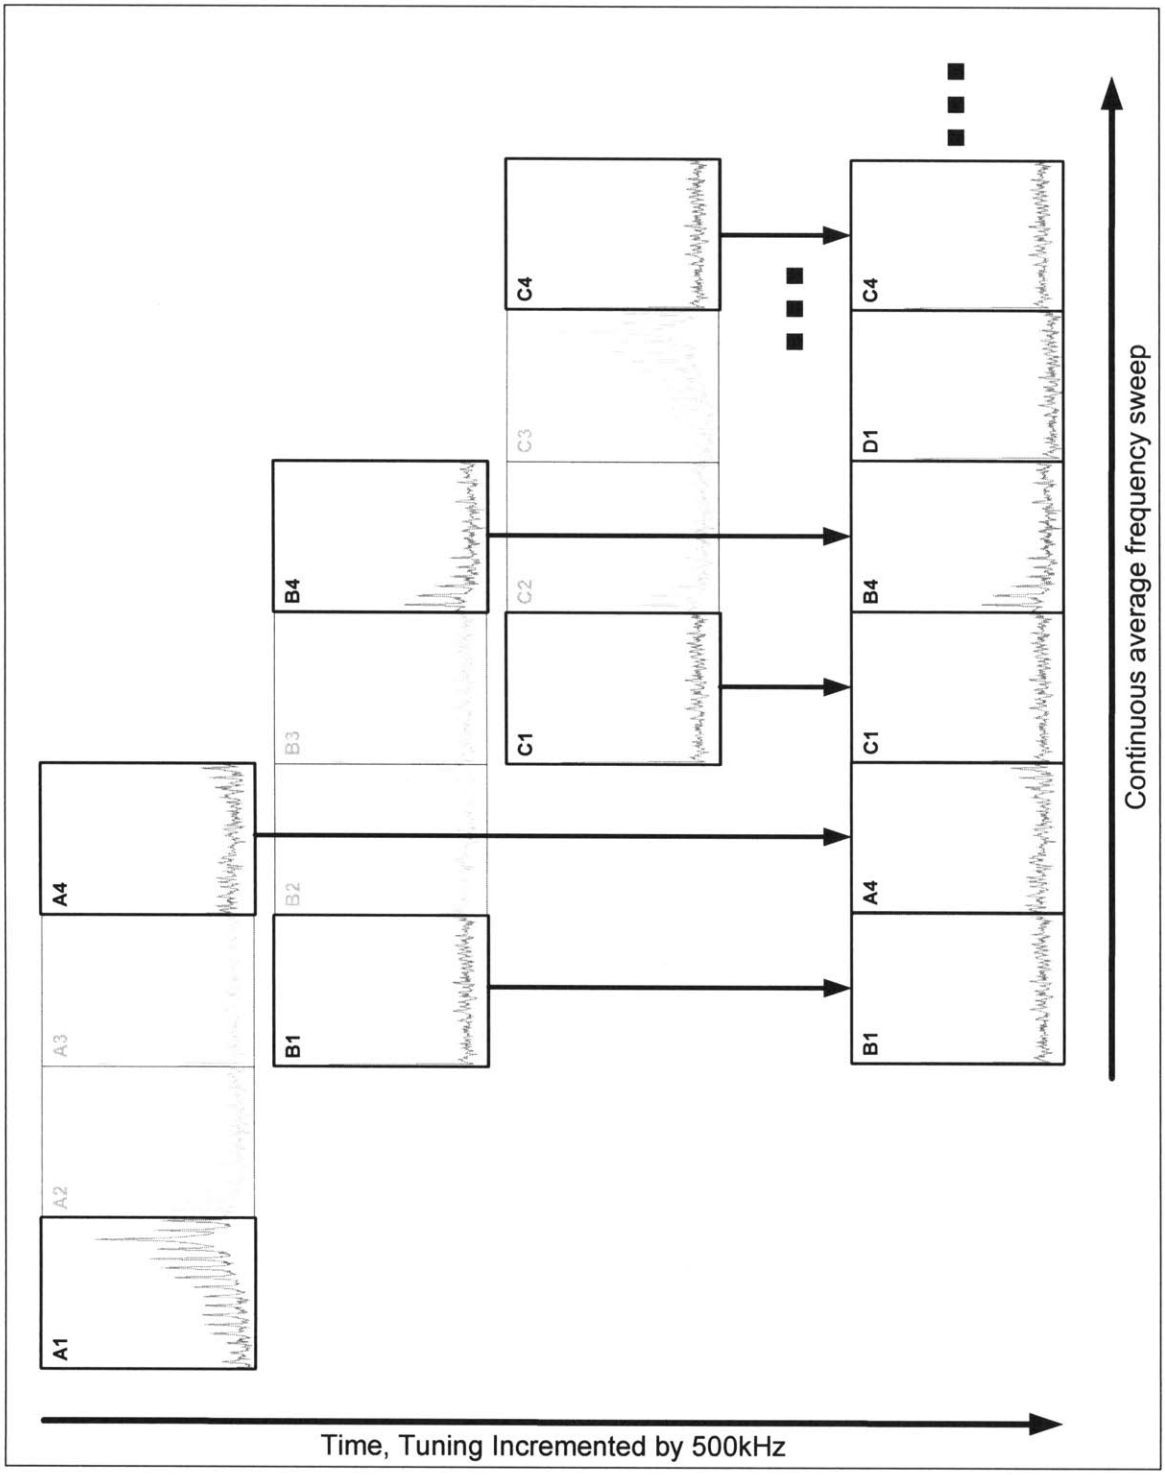

Figure 3.6 Data Interleaving Used to Perform Spectrum Sweep

As the data is recorded, a single "unroll" file is produced which contains a listing of all frequencies examined. For each frequency, the number of FFTs taken (this depends

upon the amount of delay between incrementing the tuning of the hardware), the minimum value, the maximum value, and average values are recorded. From this data, it is possible to build a spectrogram of the entire frequency sweep and to perform some analysis on the data. Pattern recognition on the data peaks is done for the NTSC television detection discussed later in this chapter. A diagram showing the complete GNU Radio pipeline can be seen in Figure 3.7. The GUI can be enabled or disabled during capture by passing options to the program.

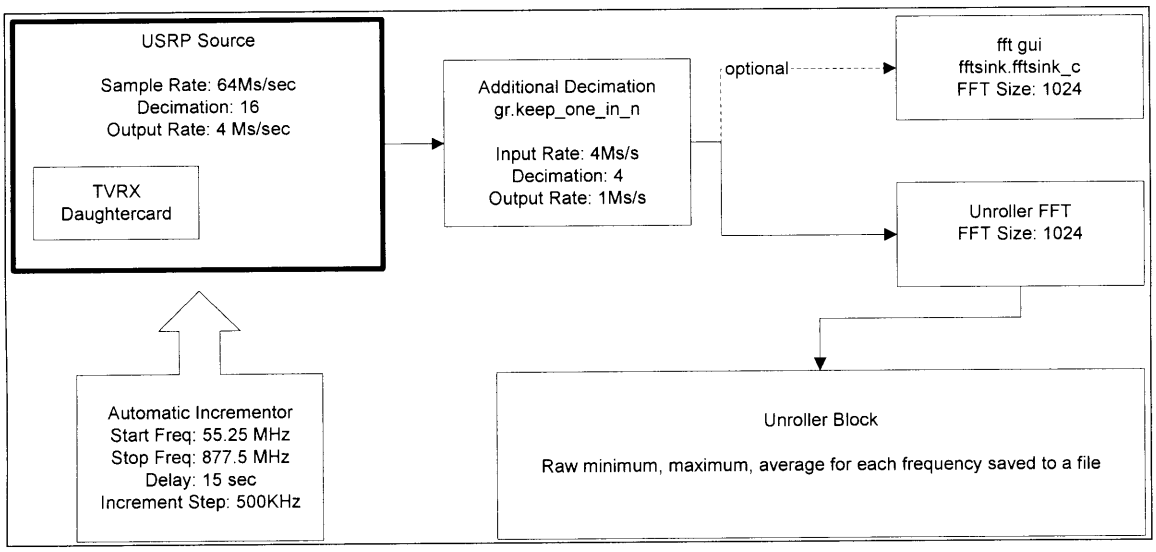

Figure 3.7 Simple Capture GNU Radio Pipeline

3.4.2 NBFM Usage Duty Cycle

A desire to study the duty cycle of NBFM radio communication inspired this module. Narrow-Band FM modulation is used for two-way half-duplex radios of the sort police and fire departments use. Frequency usage on these channels is intermittent as police officers call back and forth to headquarters. Figure 3.8 shows several NBFM channels in use around 470.3 MHz.

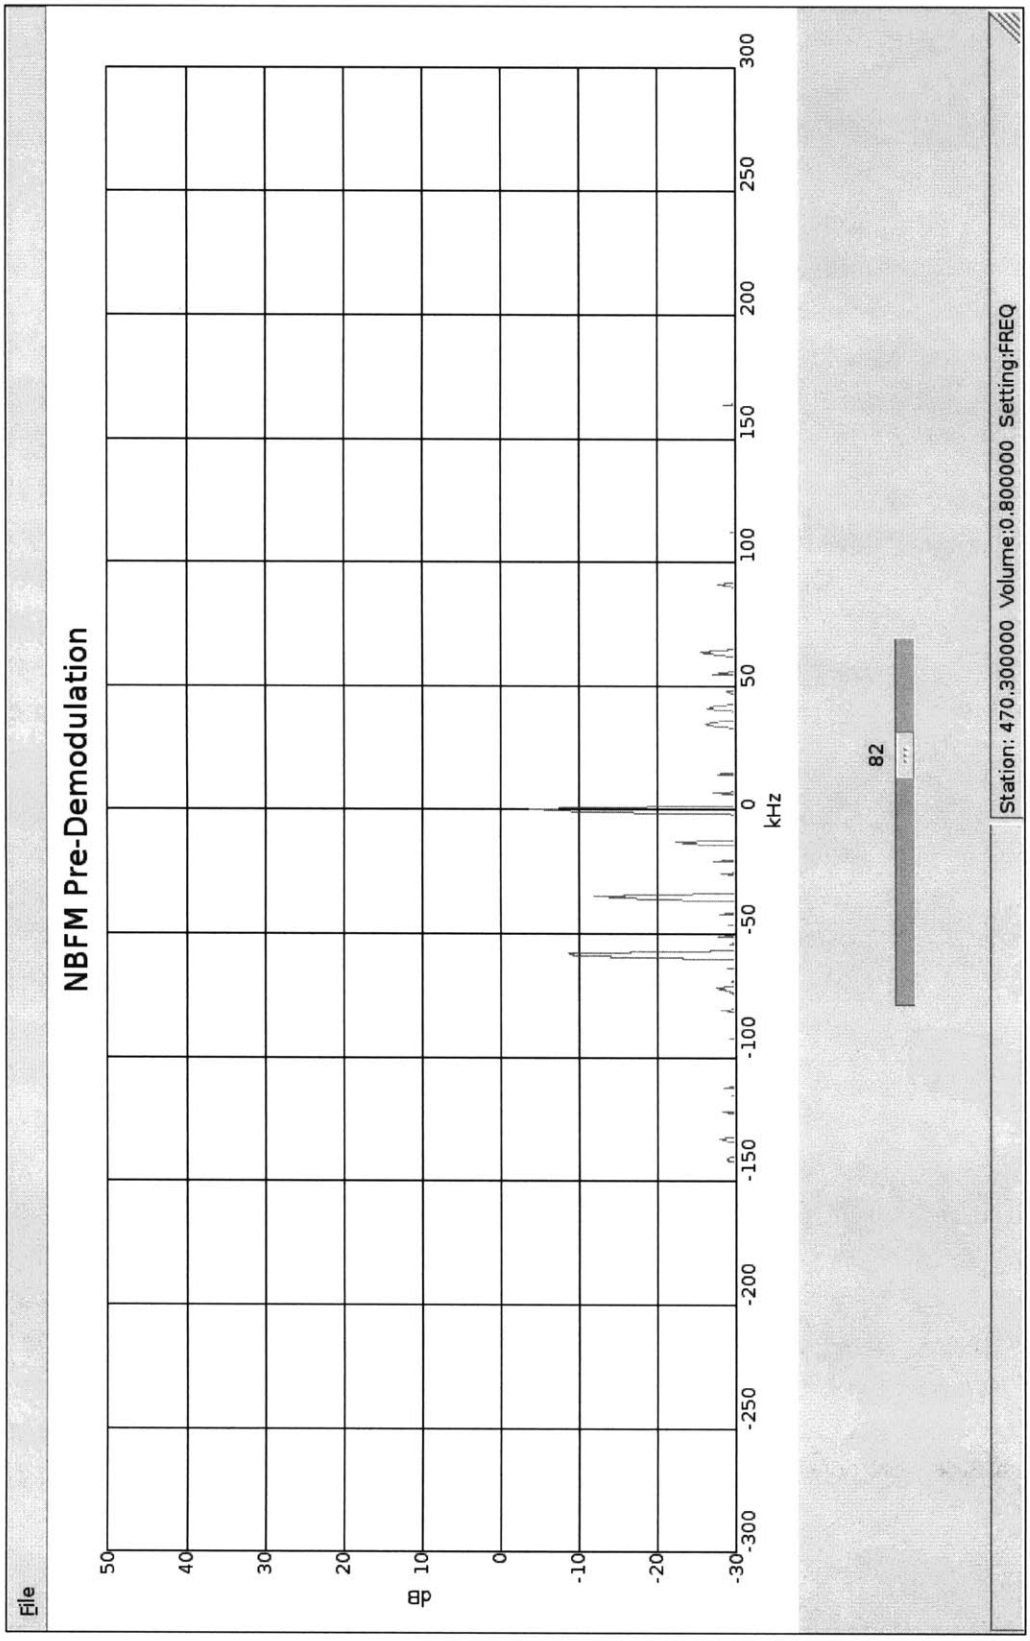

Figure 3.8 NBFM Channel Usage is Highly Periodic

The original design made used of a special threshold detection core library module. The difficulty with using this was that it was only possible to study a few known channels at a time. Instead, a FFT usage block was created to record the output of a FFT taken on the incoming sample stream. Whenever a channel spike passed the threshold (i.e., when it was in use), the time was recorded. Later post-processing built a record of all channels spotted over the course of 24 hours and served as a basis for the graphs and statistics presented in the next chapter. The GNU Radio pipeline for this design is shown in Figure 3.9.

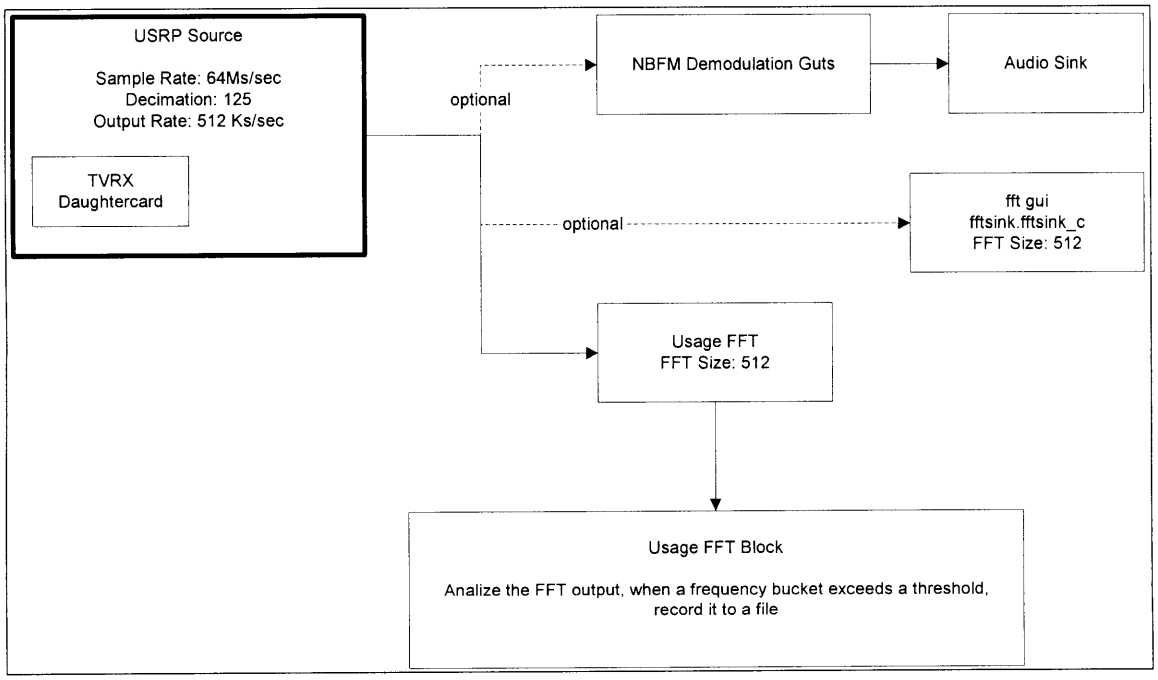

Figure 3.9 The GNU Radio Pipeline for Recording NBFM Usage Over 24 Hours

This study is the only one in the thesis to be actually carried out over a "Day in the Life" of Narrow-Band channels!

3.4.3 GPS Location Capture

The desire to correlate frequency information with physical location grew out of the initial thoughts about RF spectrum usage: namely, that it varies greatly with space. The initial study on the spatial nature of spectrum usage was attempted using the all-in-one sample capture method of the earlier architecture design.

In this effort, spectrum sweeps were performed at several fixed locations. As it turned out, the data collected was flawed in that much had been lost due to the PC's inability to maintain the high data rates needed both to read from the USRP and write to an external USB disk. Instead, this GPS Location Capture suite was developed. The idea is to capture a small region of spectrum space and do a spatial study of that limited region rather than try to capture large regions over large areas.

The GNU Radio pipeline used for data capture is shown in Figure 3.10. The architecture is similar to the Simple Capture suite, except that over time, the tuned frequency does not change. Instead, an external Garmin GPS unit is queried at fixed intervals using the PyGarmin[28] libraries and the location and data from each frequency in the FFT is recorded. Like the Simple Capture Suite, the number of FFTs, minimum FFT value, maximum FFT value, and average FFT value are recorded for each frequency.

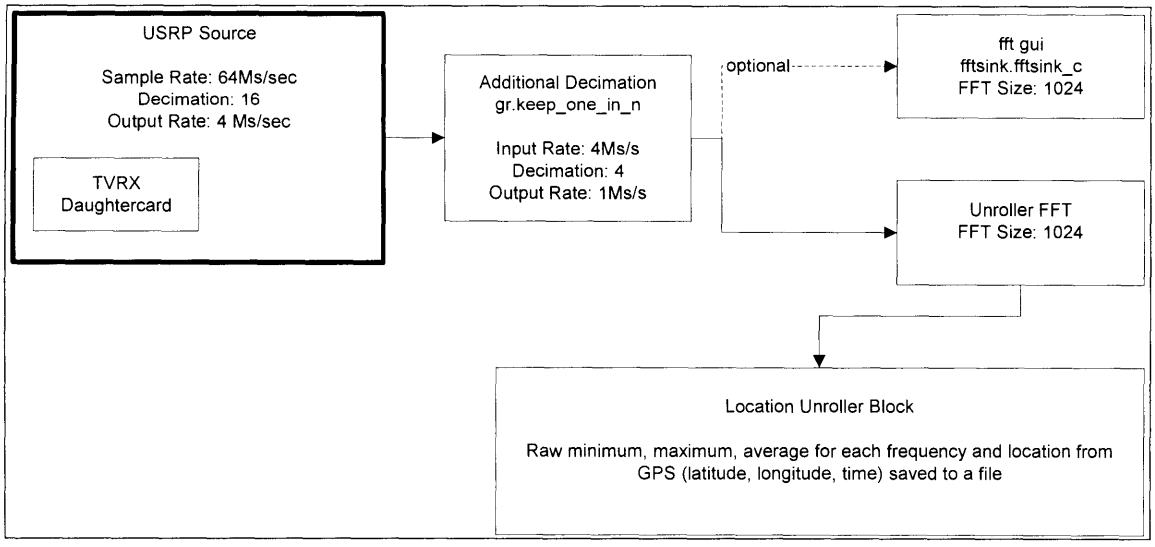

Figure 3.10 The GNU Radio Pipeline for GPS Location Capture

After the data is recorded, there are some utility scripts which post-process the data in order to sort it, consolidate frequency data from same locations, and produce an output file sorted by frequency that has values for each location. Because I used this capture to exhibit the signal power variation of MIT's FM Radio station over Cambridge, I deliberately chose the maximum amplitude from a frequency range 75 kHz above or below the desired center frequency. In wideband FM modulation, this is the maximum deviation used to frequency-modulate the signal.

A 3D interactive display was built using the Python extensions for OpenGL[29]. This consists of a satellite image mapped to a plane on top of which a graph can be built to show captured data. Any map or satellite image may be used, so long as the coordinates of two points are known and can be correlated with the pixels of the image. A view of this display as developed for the Media Lab spring sponsor meetings is shown in Figure 3.11. The data for this is from the capture attempt discussed above. The poles shown are average amplitude at a given frequency for each of those fixed locations.

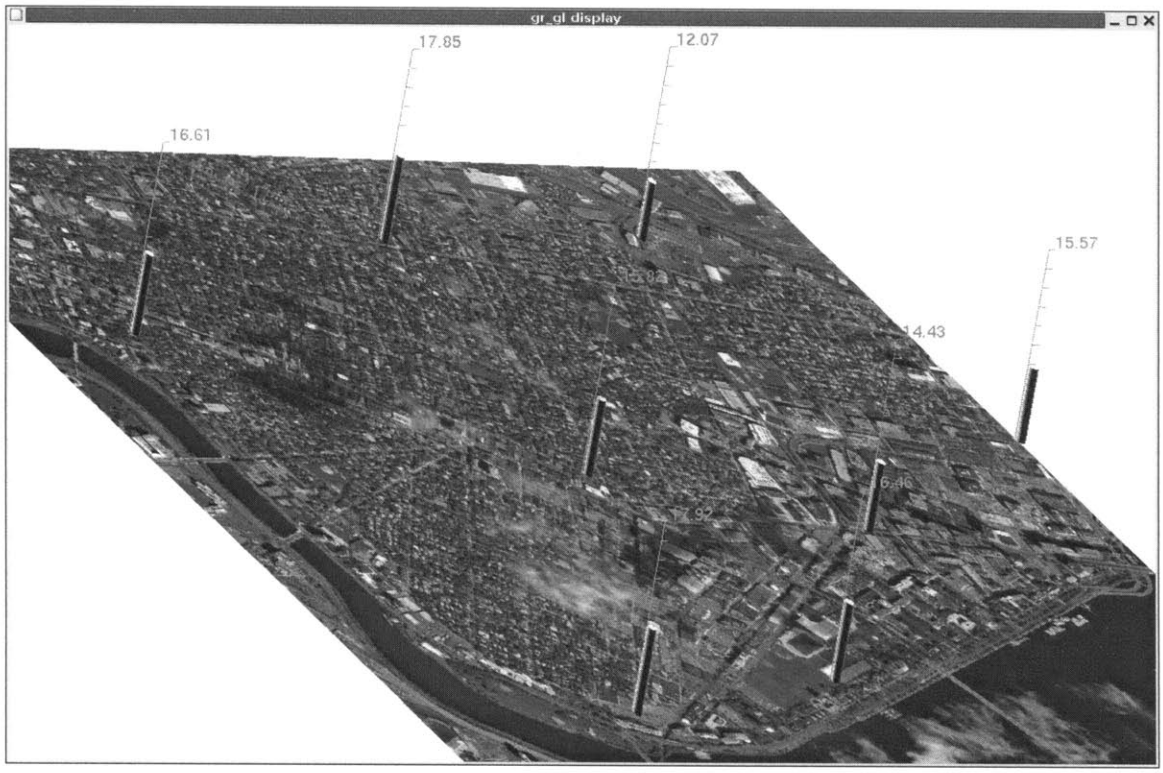

Figure 3.11 OpenGL-based 3D Interactive Display for Plotting GPS Location Data

3.4.4 FFT Pattern Analysis Using Captured Data

One of the first ideas for identifying frequency occupants in this thesis was to try to develop a "spectral signature" of a given signal. This is quite easy when analog NTSC television signals are the target of interest.

In the NTSC television standard, the chroma (color) information is broadcast 3.5795 MHz above the carrier frequency and the FM modulated sound is 4.5 MHz above the carrier[30]. By looking for frequency peaks which are these distances apart, a basic assessment as to whether a TV channel is present or not can be made. More elaborate schemes are possible, such as checking for RF energy located exactly at the scan frequencies, but this simple method is remarkably reliable. Figure 3.12 describes how such a pattern-detection might be performed on Boston WGBH Channel 2.

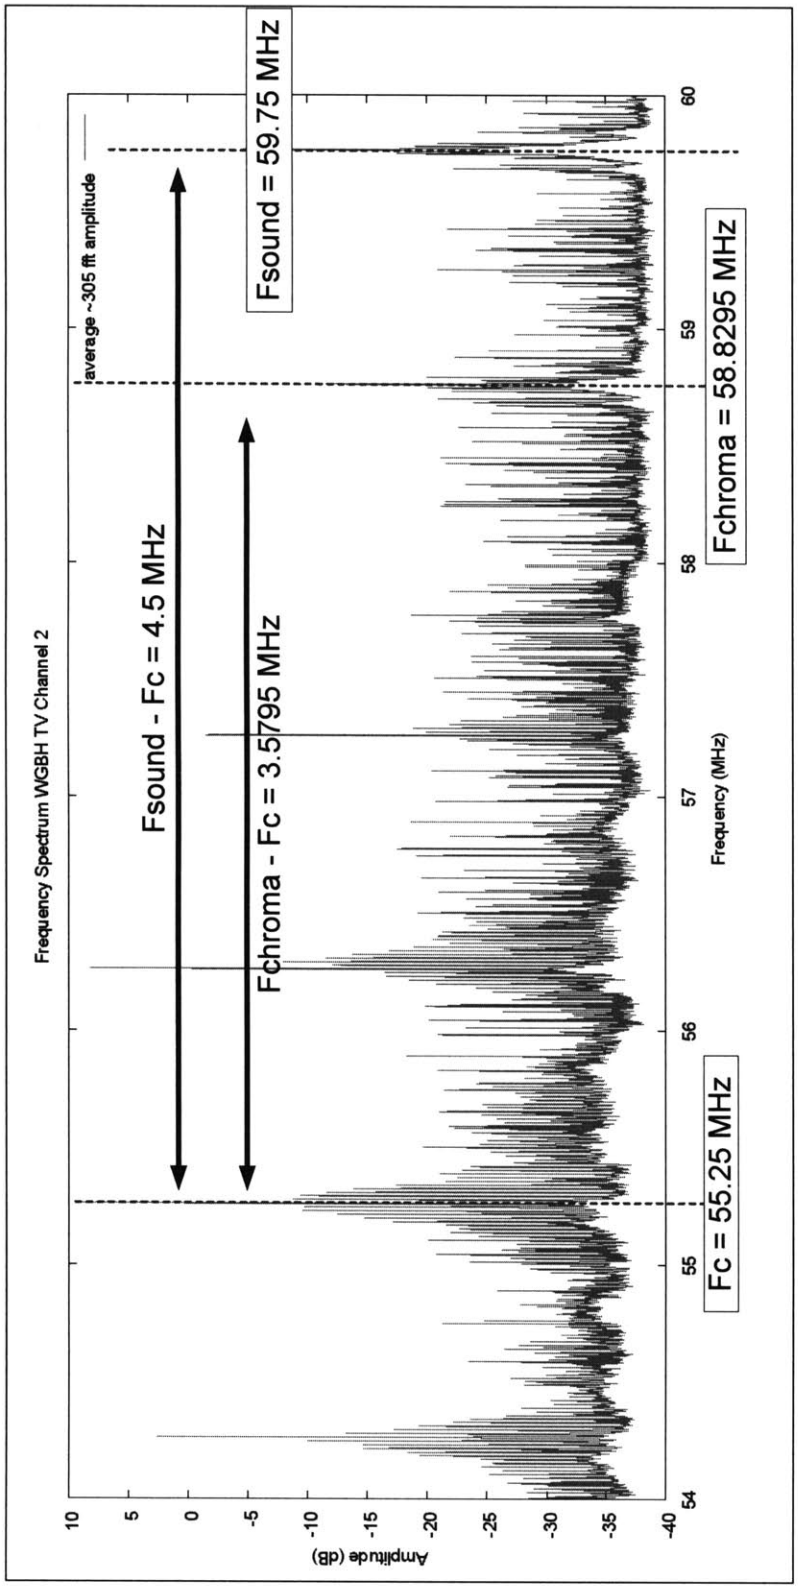

Figure 3.12 Exploiting the Signature of an NTSC Television Signal

3.4.5 Brute-Force Detection Using Demodulation

A brute-force method for detecting the nature of an RF signal is to run it through the demodulation process and asses the quality of the result. By running a signal through several different demodulations, it should be possible to determine which scheme best suits the given signal.

In the case of the code suite implemented here, two modulation schemes were applied as the capture program slowly swept through a frequency range. The number of peaks, the spread from minimum peak value to maximum, and the standard deviation of peak frequencies were recorded for both Wide-Band FM demodulation and Narrow-Band FM demodulation for each frequency. (Wide-Band is what we know as FM Radio while Narrow-Band is typically used for trunking public-service police and fire two-way radios). The signal at each station was demodulated and the output filtered to be within musical/vocal range between 0 and 4kHz. A diagram showing the GNU Radio pipeline for this design is shown in Figure 3.13.

Normally, if the proper demodulation scheme is used, it is quite clear when a station using that demodulation is selected or not. Figure 3.14 and Figure 3.15 contrast the output of this brute-force scheme using WFM modulation when a radio station is tuned in and when not. This is exactly what one who listens to FM radio is used to: when a radio station is well tuned-in, we hear music or voice clearly and when it is not, we hear static and noise. In this design, a minimalist scheme is implemented to decide whether what is tuned to is noise or music/voice. We can guess the type of modulation used for a signal based upon the demodulated output with the best quality.

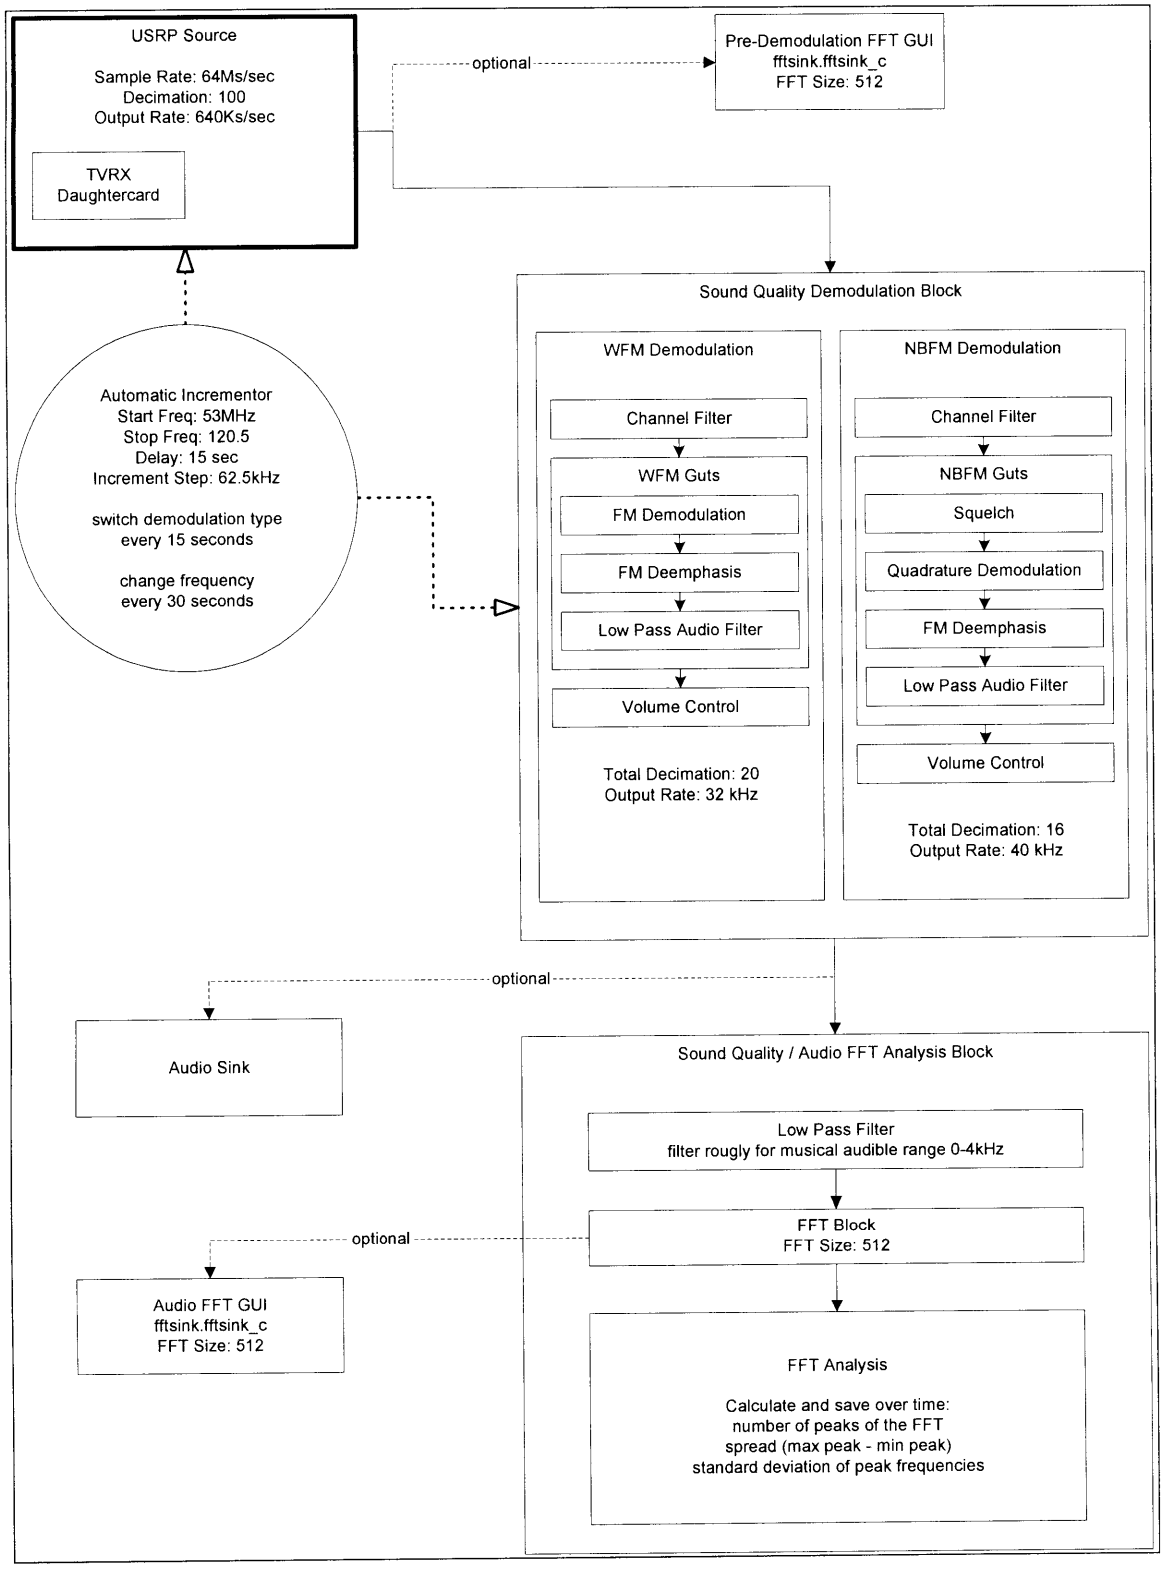

Figure 3.13 GNU Radio Pipeline for Brute-Force Demodulation Analysis

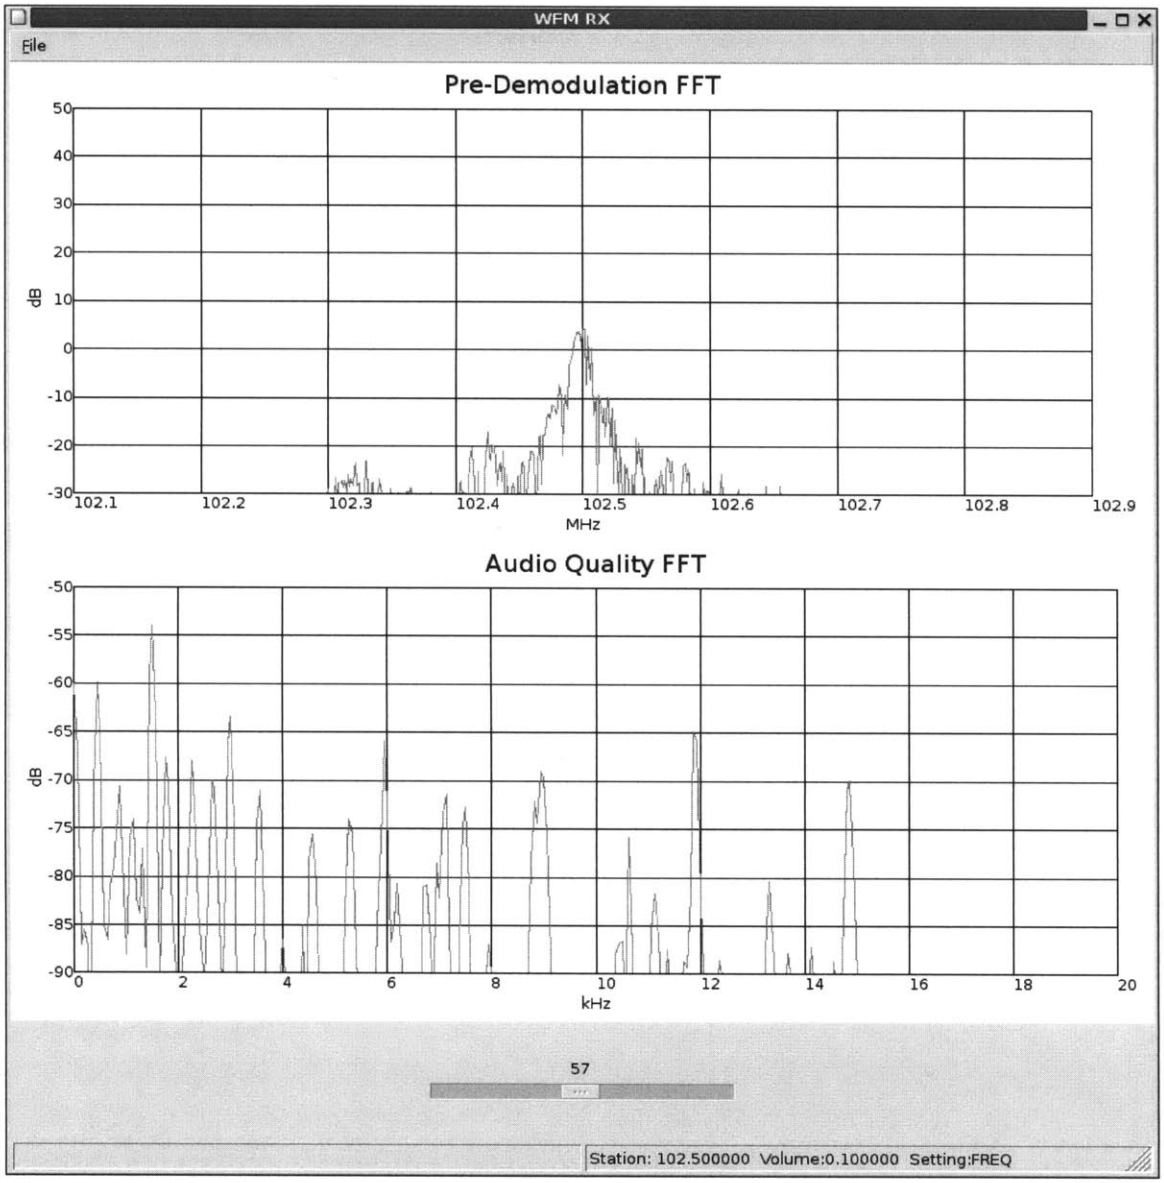

Figure 3.14 Brute-Force Identification of an FM Radio Station Using Audio Quality

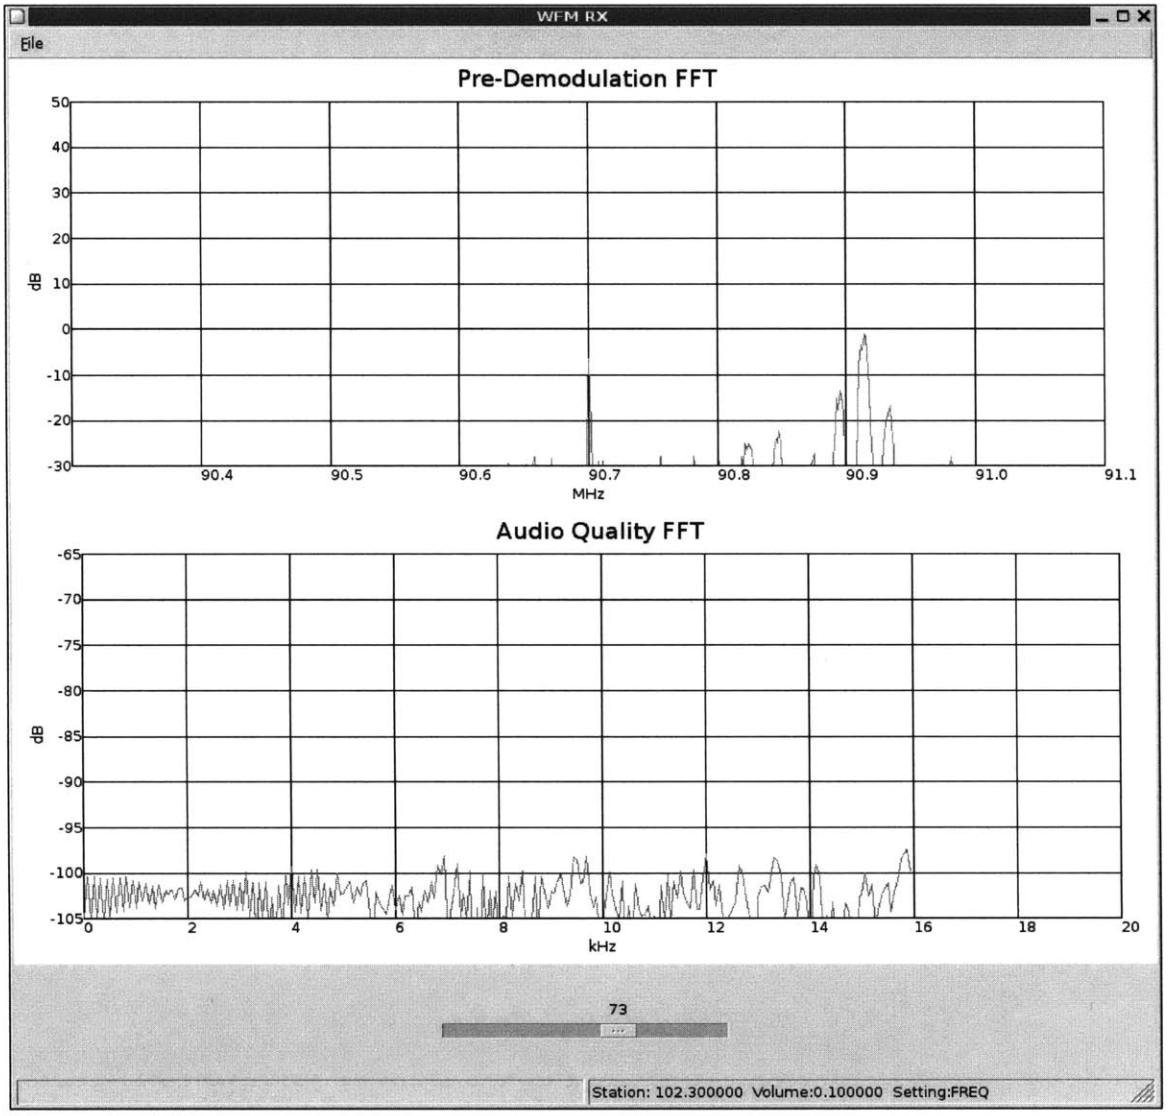

Figure 3.15 Brute-Force Identification of Noise Using Audio Quality

3.4.6 Detecting Receivers

A more-precise measure of frequency usage comes from the ability to detect receivers rather than just the presence of energy from a far-away, powerful broadcast tower. If one can say for any given moment that there is no listener of a frequency in the immediate vicinity, why can't a smart-radio system be allowed to reclaim that frequency for low-power, localized communication?

Ideally, we'd like all receivers to have some built-in transmit capability in order to announce what they are listening to. Practically, this does not exist today, but there are still ways to exploit existing radio designs in order to detect what is being listened to. The Mobiltrak Company already offers a product to track what FM radio stations people listen to in their cars as they pass on the expressway.[31] In addition, in the United Kingdom, where people must pay a television licensing fee to use a television, the government enforces the rule by driving a detector van around to detect unauthorized usage. This relies on the fact that most, if not all televisions have a local oscillator that operates at a known frequency away from the station being watched.

For this thesis, an application demonstrating local oscillator detection was built. Rather than operate across a sweep of frequencies as the other application suites do, this application is designed to detect a listener of a specific radio station. In a super-heterodyne radio receiver (as most consumer FM radios are), in order to tune to a particular station, the receiver contains a local oscillator tuned 10.7 MHz above the station being listened to. When mixed with received radio, an image of the desired station is formed at 10.7 MHz (which can then be further demodulated and down-converted for listening. In this application, two TVRX Frontend USRP daughter cards are used. The first is tuned to the FM radio frequency of interest and output to the sound card as a reference and the second is tuned 10.7 MHz above. When a nearby radio is tuned to the same frequency as the reference audio, a clear spike is visible on the display. Figure 3.16 shows the GNU Radio pipeline for this design.

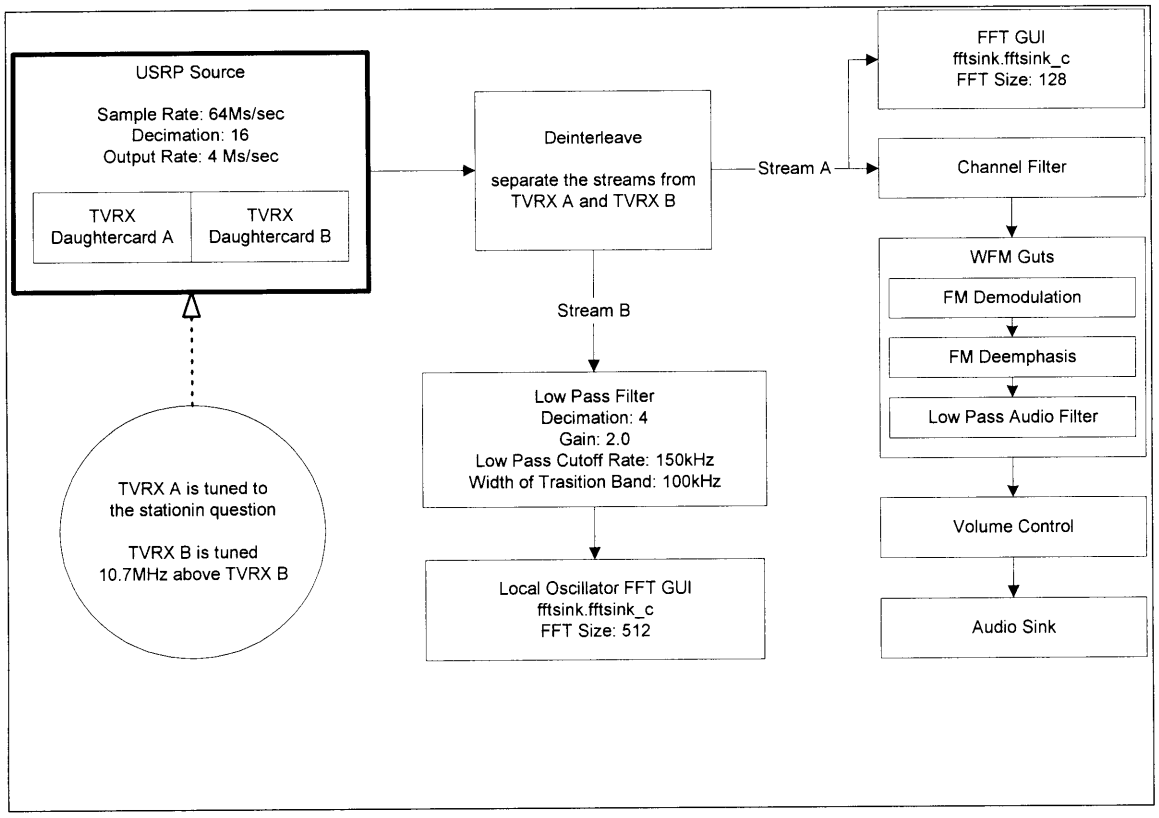

Figure 3.16 GNU Radio Pipeline for Local Oscillator Detection

An alternate design using an interferometer is presented as a design which may do a better job at localizing usage. One difficulty in successfully recognizing a local-oscillator spike is that it often falls well within the range of other radio stations.

In order to realize this design, the same type of antenna and same length of wire is attached to each TVRX daughtercard in order to make sure the impedance of each path matches as closely as possible. A grounded shield is placed in-between the two antennas. The theory of operation is that large, broadcast-tower emissions will still reach each receiver antenna in essentially the same way. Smaller, localized signals (such as the unintentional local-oscillator radiation), should be received noticeably better from one antenna or the other. When one received signal is subtracted from the other, the difference will remain.

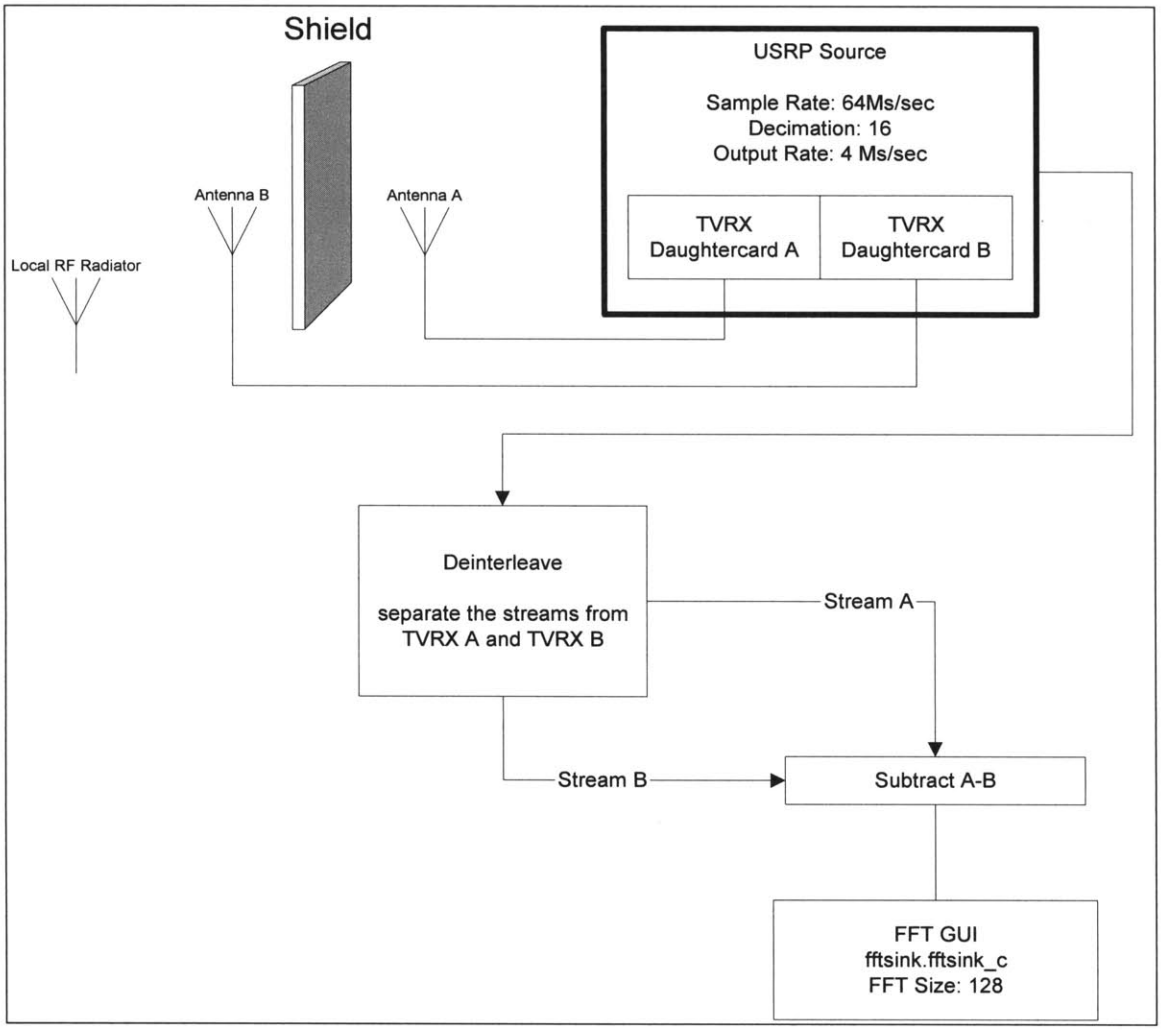

Figure 3.17 Proposed GNU Radio Pipeline and Experimental Interferometer Setup to Better Detect Local Oscillator

Although not implemented for this thesis, this design should be an improvement that can be built for future study.

4 Results

This chapter describes results of experiments carried out using the GNU Radio modules designed and built for this thesis as described in the previous chapter.

4.1 Basic Energy Map of Spectrum Usage

Using the Simple Sweep module, a basic energy map of spectrum usage from about 50 MHz - 860 MHz in the Cambridge area was generated. Although this does not represent a complete picture of RF Spectrum usage in this range, it does convey a general sense of occupied bands.

This map is shown in Figure 4.1. The most prominent peaks on this graph are FM radio stations around 100 MHz and television from 500 – 650MHz.

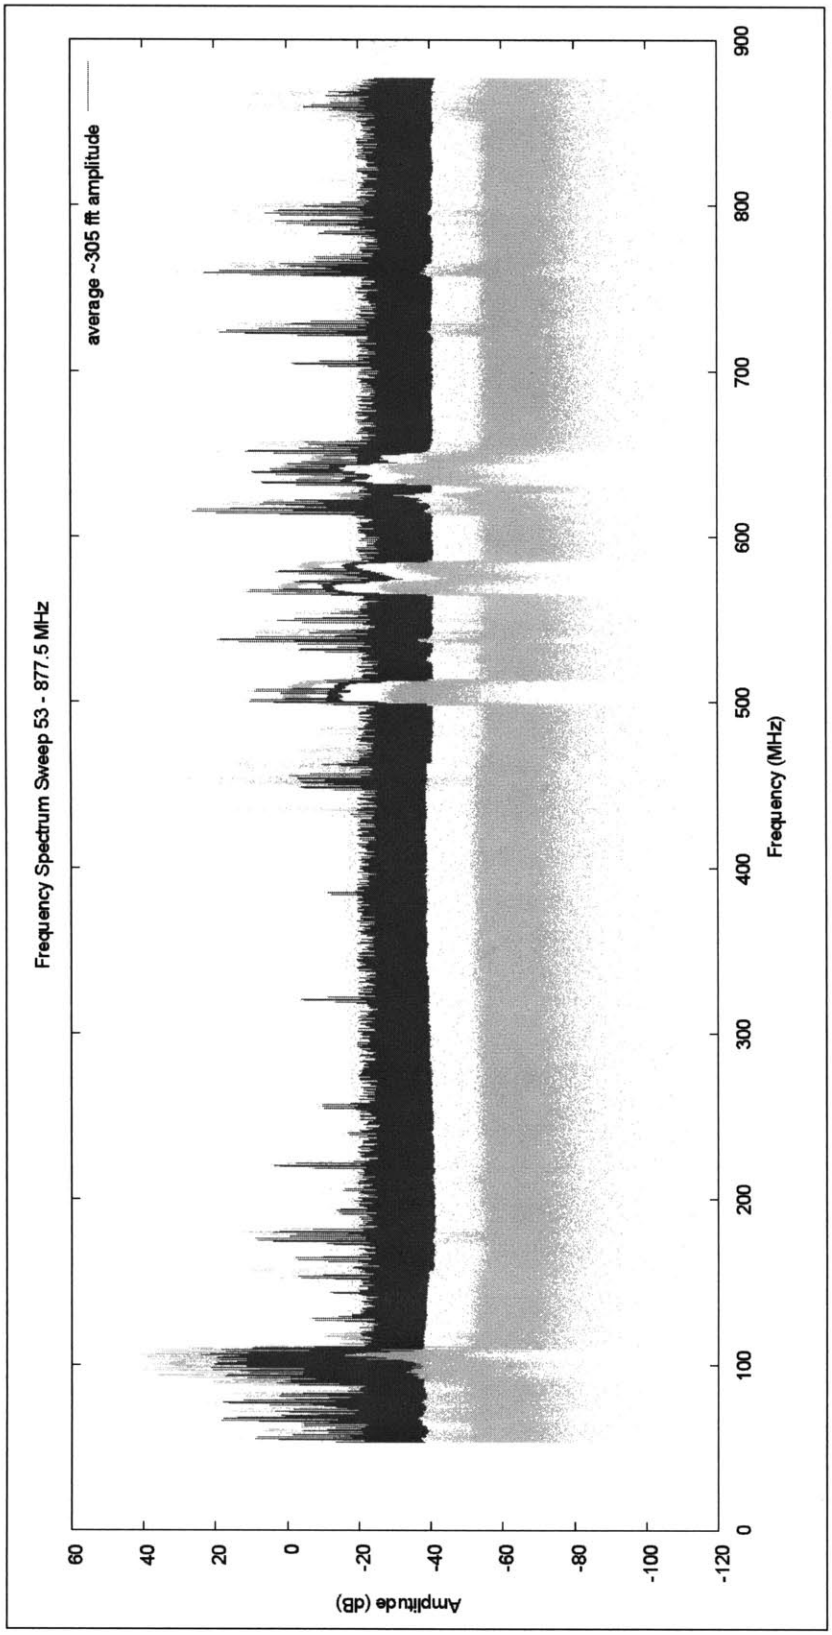

Figure 4.1 Sweep of RF Spectrum Space From 53 MHz– 877.5 MHz.

4.2 Usage Over Time of Highly-Periodic Signals

In order to better understand and illustrate how some frequency usage varies highly with time, a study of Narrow-Band FM signals was performed over a 24-hour period from midnight to midnight on the next day. A summer weekend and a range that included known Cambridge and Boston police frequencies was deliberately selected in order to study what is presumed to be a maximum usage case on a typical day. Only during emergencies or holidays, might one expect to find even busier usage. Over the course of a full day, channel usage was recorded and is presented as the histogram in Figure 4.2.

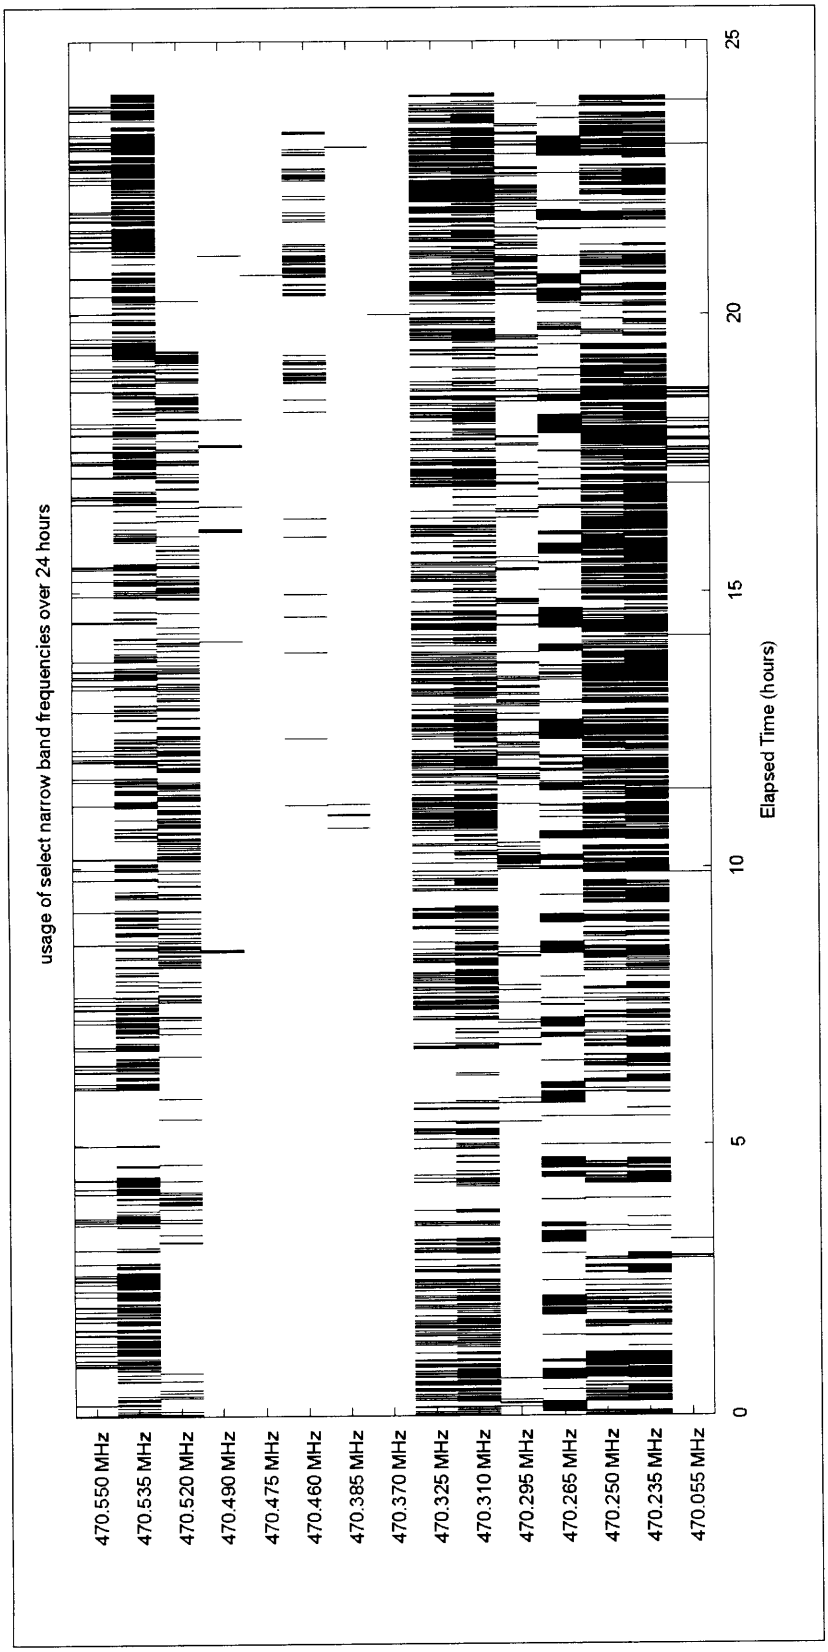

Figure 4.2 Select Narrow-Band FM Usage Over A 24-Hour Period

Usage across these bands varies greatly. The frequencies from 470.370 MHz – 470.520 MHz are clearly not used very much and for the most part, just during the daytime hours. 470.295 MHz and 460.055 MHz are also seldom used in the middle of the night. The other frequencies shown are mostly police and fire communication systems and show heavy usage for most of the 24-hour period, with slightly less density in the early morning hours.

While one might easily conclude from this that public-service frequency usage is constant, this graph is misleading in showing high-density constant usage. This was surprising, because it was assumed that even the most busy frequencies would not have a very high duty cycle. Clearly, the lesser-used frequencies are used far less than the densely-used ones. Data was recorded as fast as possible, so any small amount of usage lasting as little as a few seconds would appear at this scale as a dark bar and help create the illusion of constant usage.

In fact, even the frequencies used the most are actually used a small percentage of the total time over 24-hours, the highest-used frequencies used less than 20% of the time. Figure 4.3 shows the overall duty cycle percentage using the same data.

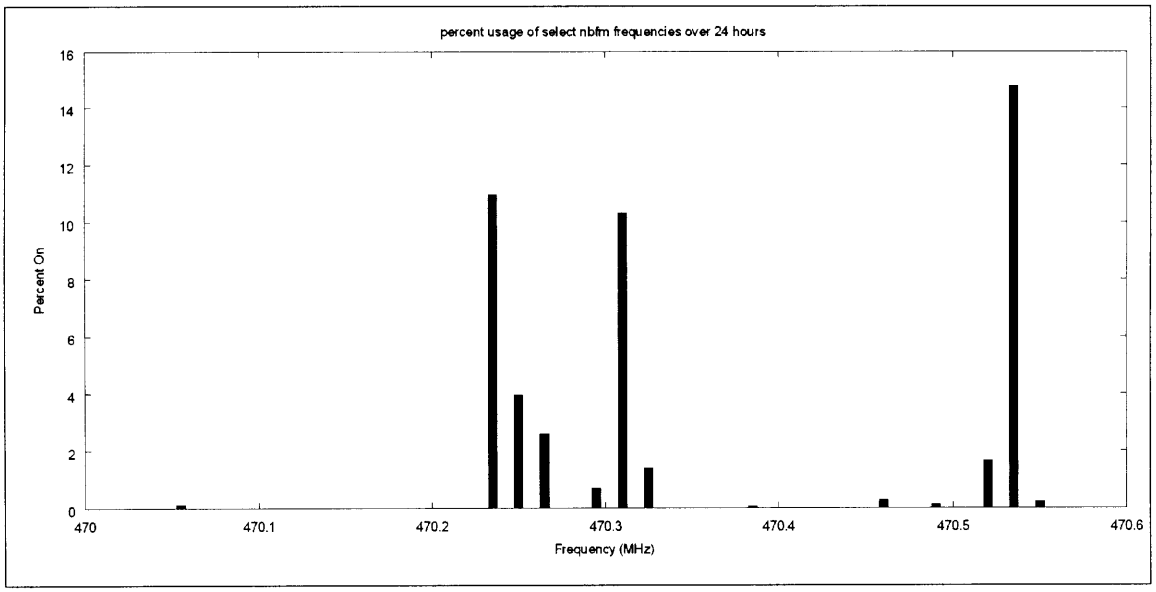

Figure 4.3 Select Narrow-Band FM Percent Duty Cycle Over A 24-Hour Period

With usage less than %20 during what is presumed to be a busier day on busier channels, clearly, these channels are an example of inefficient use of spectrum space. The mere suggestion of re-allocating or reusing channels dedicated to emergency personnel is understandably suggested with great caution. It would be possible to replace existing antiquated systems with more-modern digital ones that could use the space more efficiently. There are also other design possibilities that could guarantee the right of emergency response personnel to trump other usage when needed.

4.3 Variance of Signal Power Over Geography

It is no secret that radio signals vary greatly with geography. But, the variance is not as simple as theoretical fading models. Buildings, vehicles, and other obstructions all contribute to the radio landscape; sometimes reception is improved and sometimes it is not.

For this study, MIT's FM Radio Station, WMBR 88.1 was received by the USRP system from a moving vehicle and plotted on a customized three-dimensional map of Cambridge. Care was taken to select the maximum power amplitude within the modulation range in order to compare like values over the landscape as described in the previous chapter. An overhead view of the drive-test route is shown in Figure 4.4. The transmitter for 88.1 FM is located on the top of MIT's Eastgate building in Kendall Square and is illustrated by an arrow in this figure.

Figure 4.4: Map of Drive-Test Track Through Cambridge for Signal Strength Variance Study.

Figure 4.5 is a three-dimensional view of the maximum amplitudes detected in the physical environment. The broadcast tower is represented as a tall bar at right. Figure 4.6 plots the radial fading characteristics as one is further away from the transmitter along the drive-test route. A highly-complex radio propagation environment is immediately obvious.

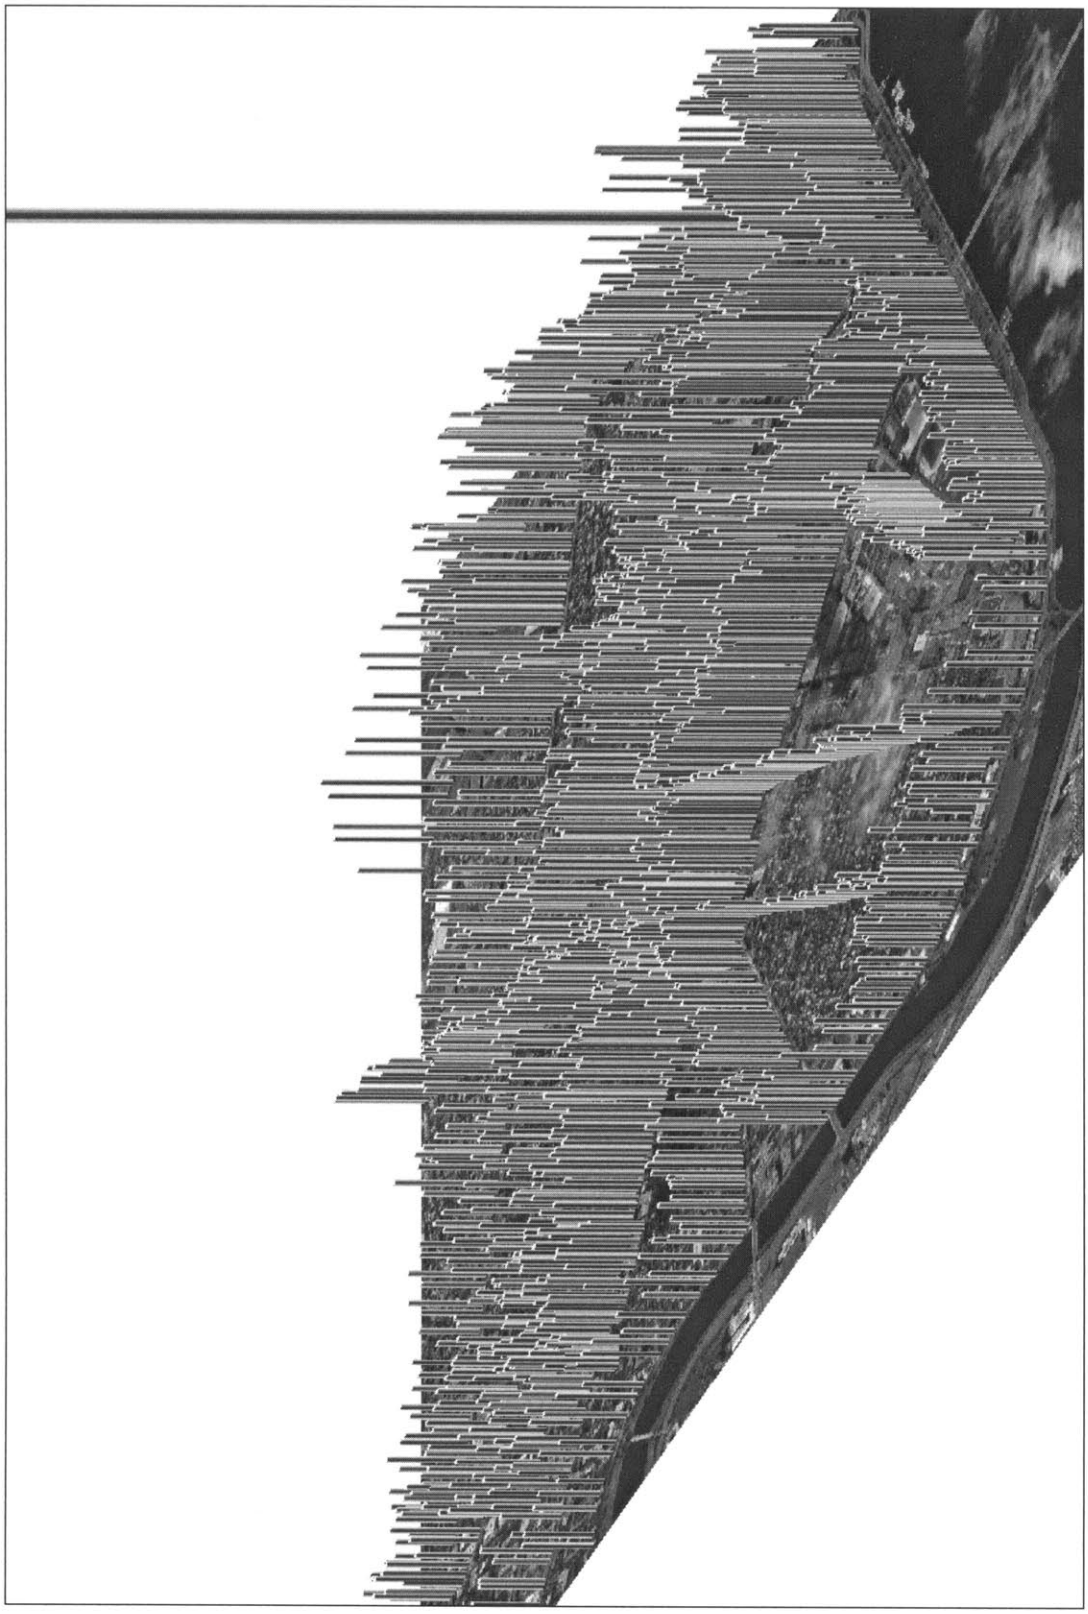

Figure 4.5 Three-Dimensional Map of Relative Maximum Signal Strengths of 88.1 FM.

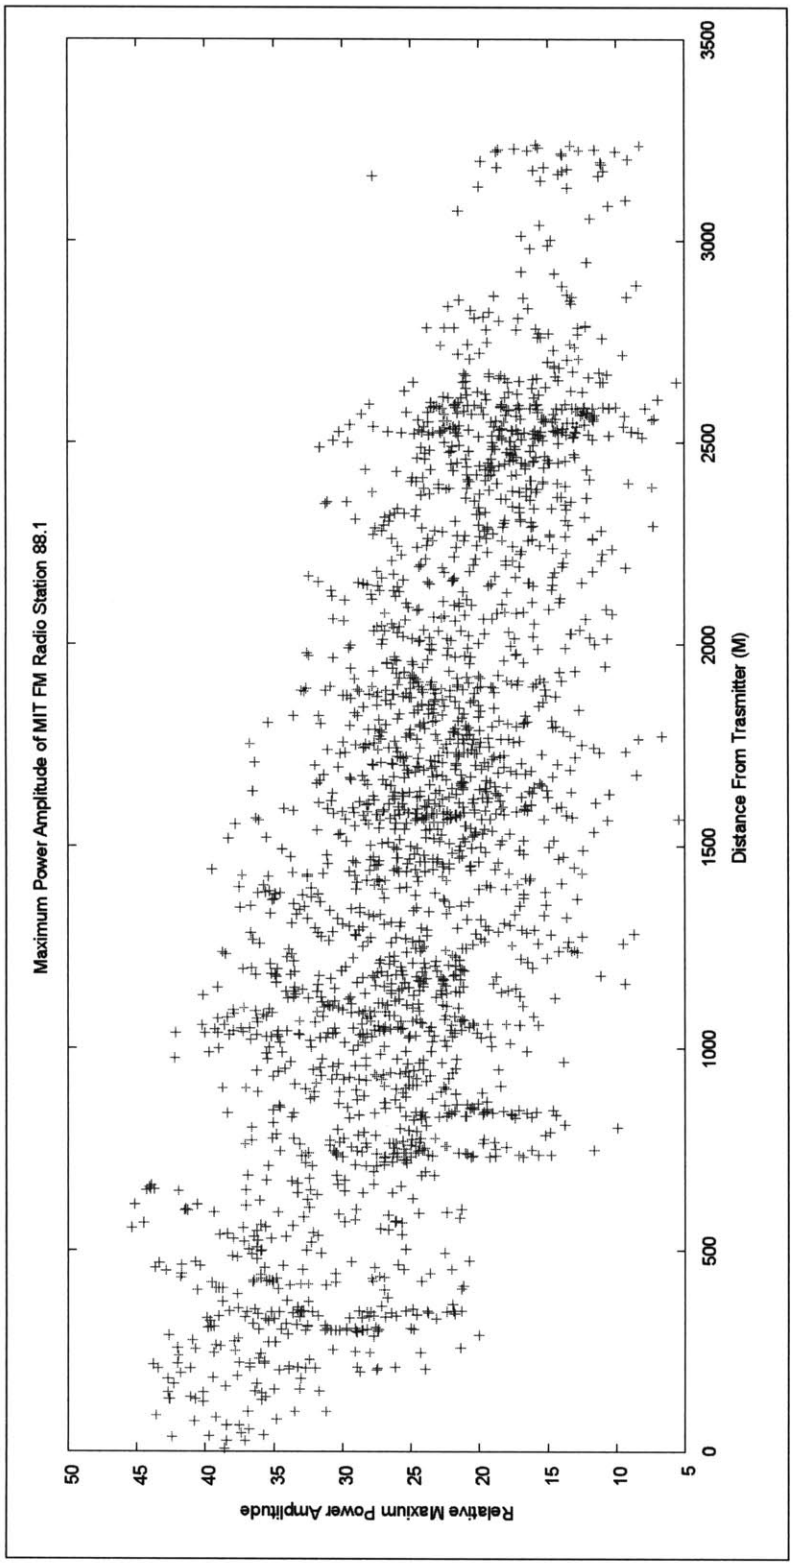

Figure 4.6 Fading Plot As a Function of Distance From the FM Broadcast Tower

4.4 Frequency Peak Detection and Pattern Analysis

One of the first ideas to emerge when developing techniques to qualify spectrum space usage was to search for readily-identifiable patterns in the spectrograms of signals. An easy choice to try this technique was NTSC television broadcasts. The details on how this detection works is described in the previous chapter, but in basic terms, the modulation scheme introduces well-recognizable peaks that are known distances from each other. By measuring the distances between these peaks, one can guess whether a frequency range contains an NTSC television signal.

The data captured using the Simple Sweep suite of software was run through such a detection algorithm and the results are presented in Figure 4.7. The algorithm results as well as actual Boston television stations are listed for comparison. Although the simple algorithm implemented for this thesis only relies on three known peaks for detection, the pattern is unique enough to be a fairly clear indication of station presence.

Two observations about the algorithm success can be made. First, the 6 MHz bandwidth of a television station was identified multiple times for several of the stations. These are shown as stacked bars on the figure. The duplicate identification is probably because in addition to identifying the highest peaks of the main carrier frequency, chroma, and sound carriers, there are smaller peaks located related to the horizontal and vertical scan frequencies (these peaks are visible in the figure of the previous chapter) which also meet the distance criteria. Second, several mis-identifications were made. For instance, the algorithm mistakenly identified a television station around 100 MHz, a location known to be used for FM Radio broadcasts.

The accuracy can be improved by making the pattern-detection criteria more sophisticated. Perhaps, in addition to detecting the maximum peaks of the baseband carrier, chroma carrier, and sound carrier frequencies, a search for the lesser scan frequencies which also occur at known offsets could be used.

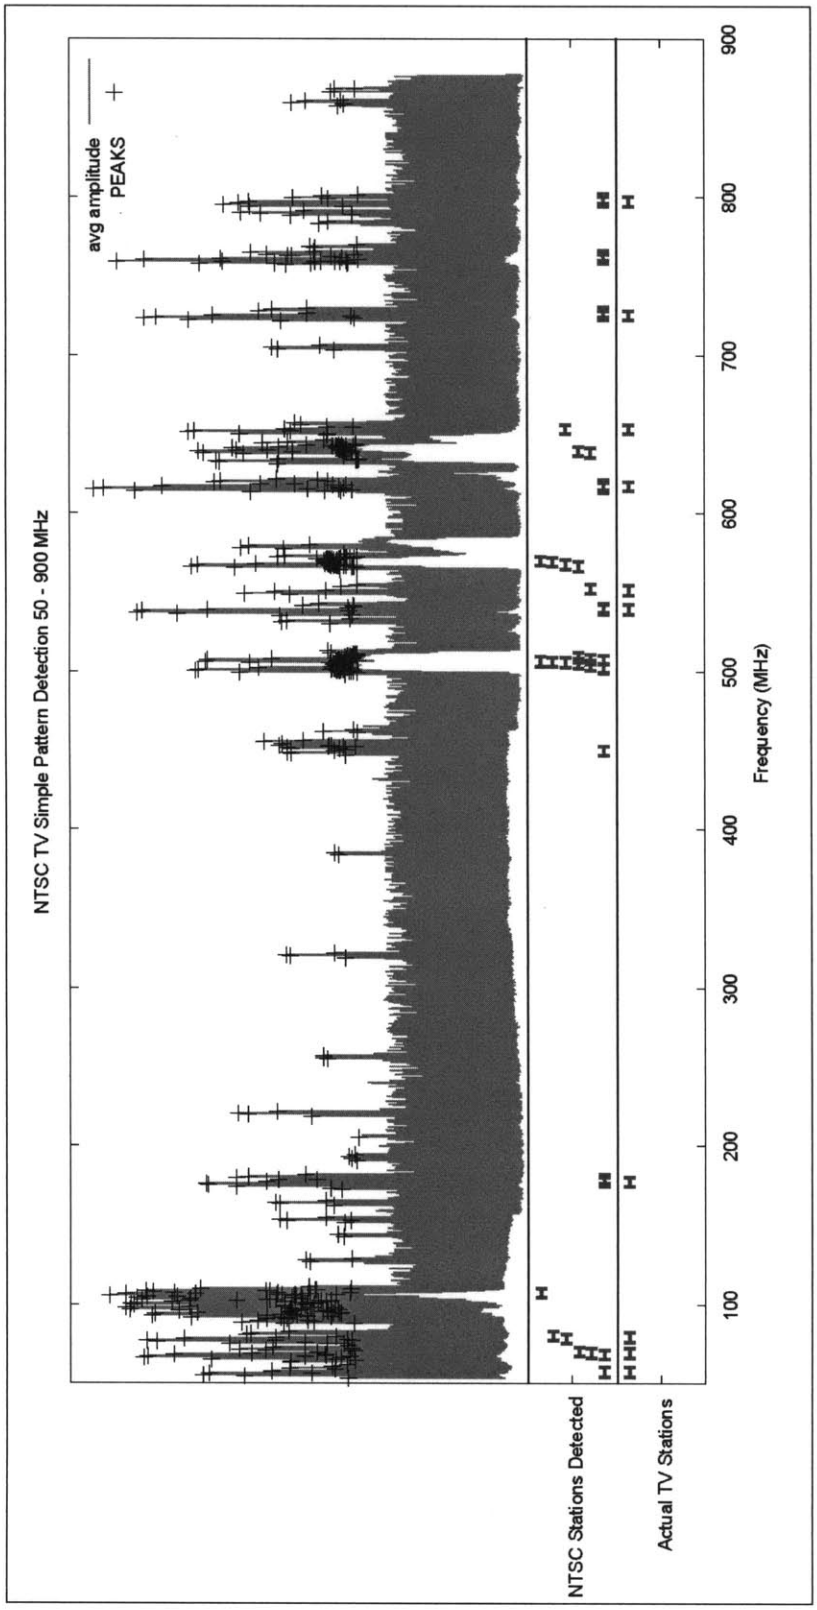

Figure 4.7 NTSC Frequency Pattern Detection Algorithm to Identify the Presence of Television Stations

4.5 Brute-force Demodulation Quality Analysis

This section describes the results of a brute-force analysis technique applied to identify FM radio stations. In this technique, a given signal is passed through two different demodulation blocks and the demodulation result with the best quality is assumed to be the proper one. Three metrics for studying the "quality" of a demodulated signal were used: the number of peaks, the spread from maximum amplitude to minimum, and the standard deviation of peak frequencies.

4.5.1 Quality Analysis Using the Number of Peaks

The peak detection algorithm used detects a finite number of peaks located no more than a specified distance apart from each other. The algorithm also requires peak frequencies to meet a minimum threshold to be considered as peak candidates. As illustrated in the previous chapter, when a station is well-tuned in using the proper demodulation type, music and/or voice frequencies appear as distinct peaks. When a station is not tuned-in or the wrong type of demodulation is used, the "audio" output is noise and is characterized by many, low-amplitude constantly-changing "peaks." The result of applying this algorithm to a range covering FM radio is shown using WFM demodulation and NBFM demodulation in Figure 4.8 and Figure 4.9 respectively. Actual FM radio stations are shown in both graphs as vertical stripes.

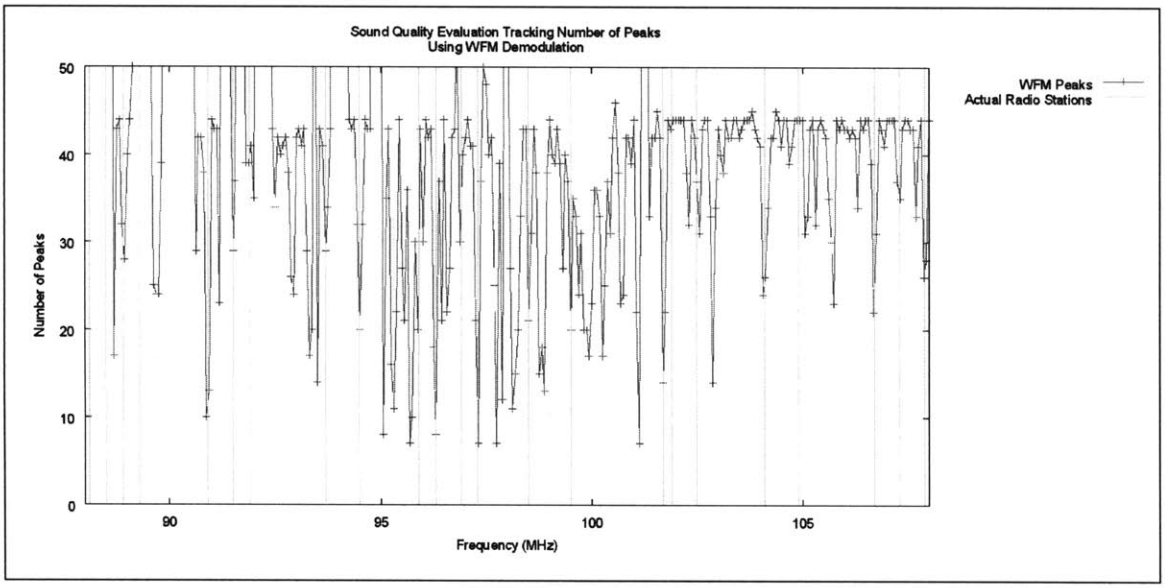

Figure 4.8 Brute-Force Demodulation Quality Using WFM and Number of Peaks.

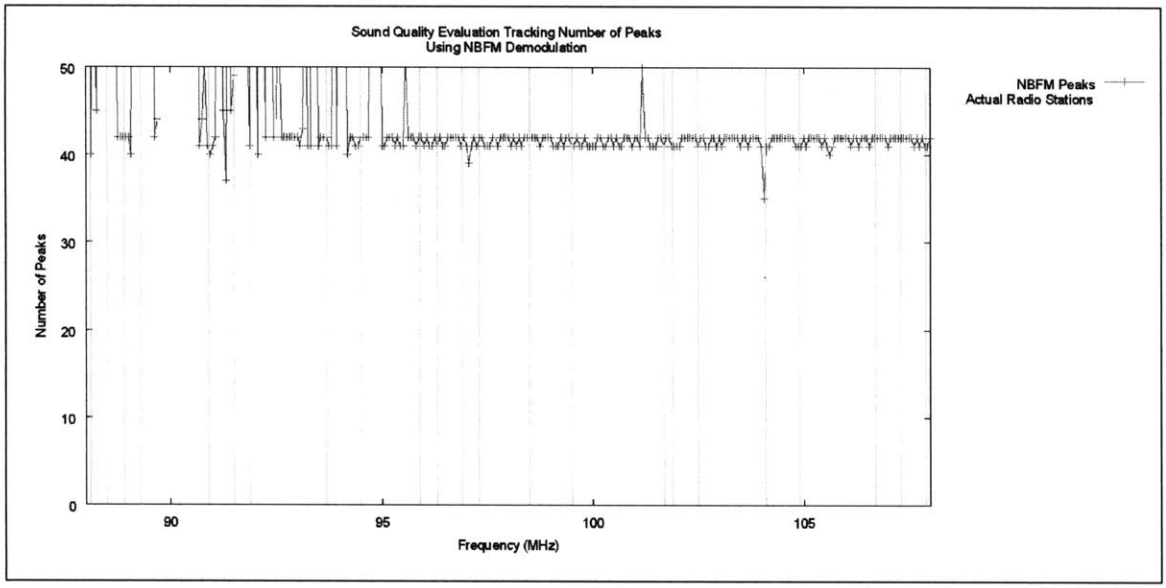

Figure 4.9 Brute-Force Demodulation Quality Using NBFM and Number of Peaks.

When using this metric, a lower number of peaks identified means that a given frequency is more likely to be an FM radio station when the WDM demodulation is applied. The algorithm performs fairly well, it is assumed that the general variance of successful identifications between 10 and 25 peaks is due to the complexity of music/voice signal present. More peaks could mean more musical harmonics present as

in a piece of music with many instruments. Given the parameters of the peak detection algorithm, 40 peaks indicates noisy signal as in the out-of-tune or wrong demodulation (NBFM) cases. A value well above 40 (above the range of the graph) was artificially set in order to indicate that no noisy peaks even met the peak-detection algorithm's minimum peak threshold. There are a few cases in both graphs where this is the case.

Worth mentioning is that this algorithm also correctly identifies the FM modulated audio carriers of low-VHF television stations as shown in Figure 4.10.

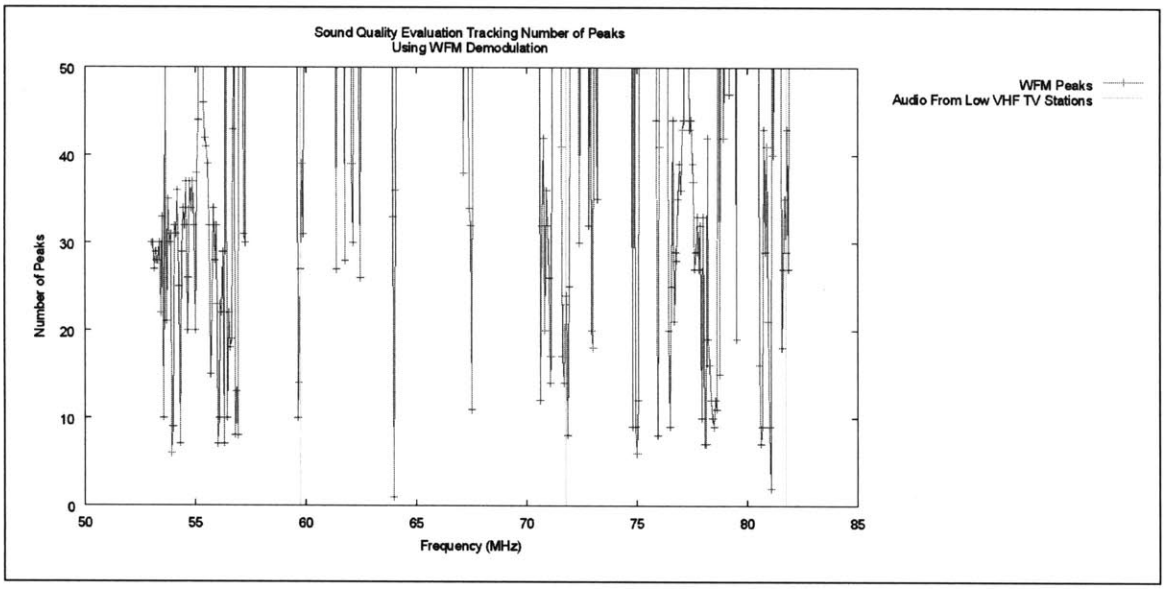

Figure 4.10 Brute-Force Demodulation Quality Using WFM and Number of Peaks Applied to TV Audio Carriers

4.5.2 Quality Analysis Using Average Amplitude Spread

The average amplitude spread is a measure of the amplitude difference from the lowest to the highest value, averaged over the frequency range of interest. Also with the intent of identifying noise from real audio content, one would expect that with real content, the spread would be greater than that without. This too is evident in the example illustration of the previous chapter. A well-tuned station has distinct peaks while an ill-

tuned one, noise is characterized by no major peaks and all amplitudes relatively the same. The results of this analysis using WFM demodulation followed by NBFM demodulation applied to the FM Radio range is shown in Figure 4.11 and Figure 4.12 respectively. Actual FM Radio stations are shown as vertical stripes in both graphs.

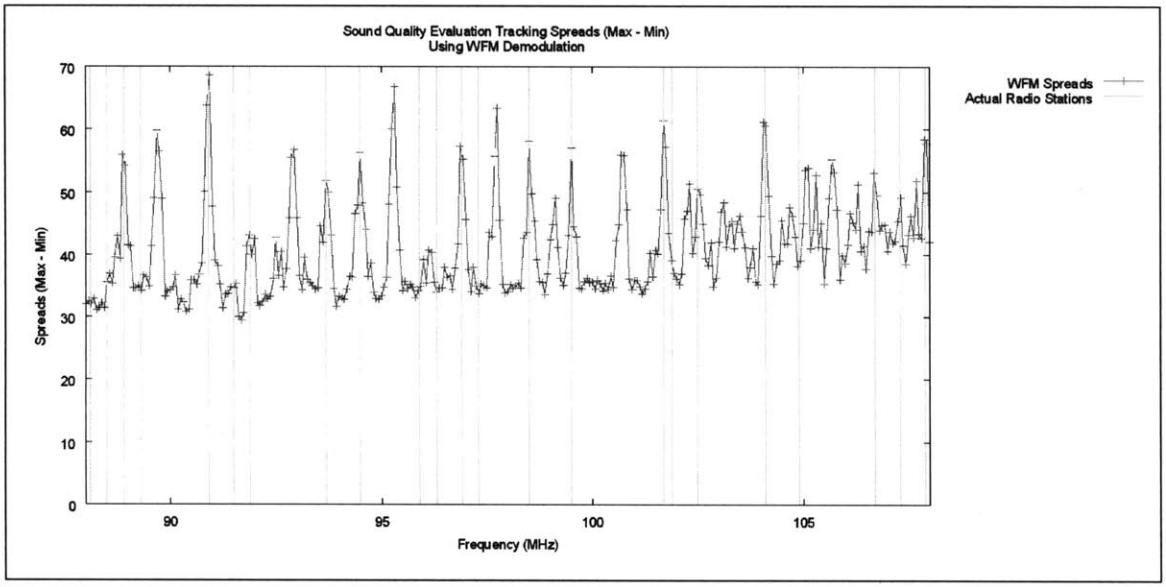

Figure 4.11 Brute-Force Demodulation Quality Using WFM and Spreads

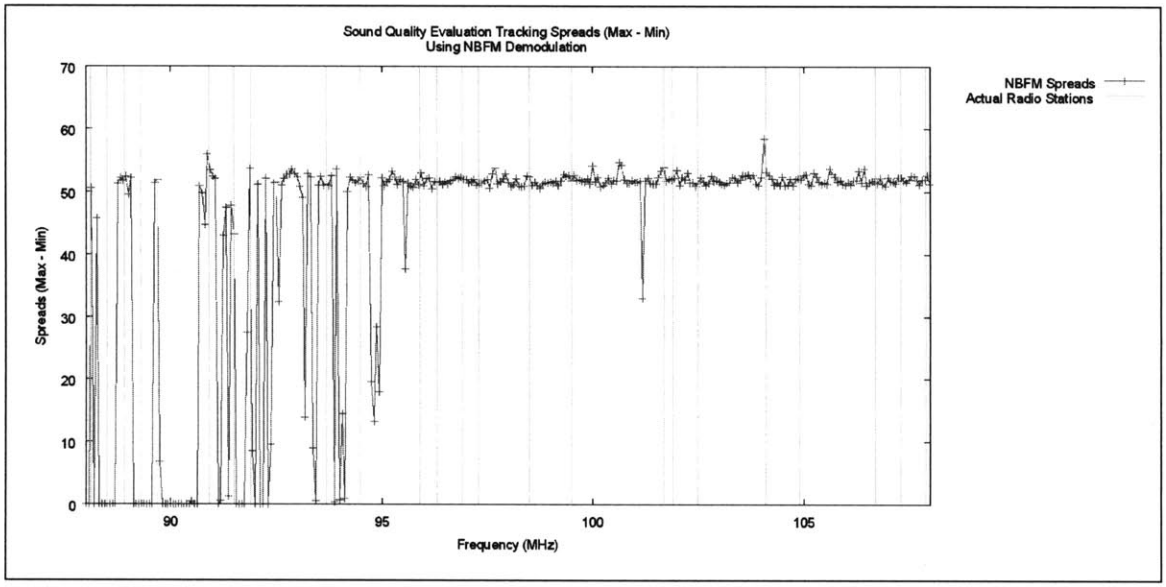

Figure 4.12 Brute-Force Demodulation Quality Using NBFM and Spreads

Using the spreads also seems to be a good measure of the audio quality. For the most part, the peaks fall on the actual FM radio stations in the WFM graph. In the NBFM graph there are few peaks present.

4.5.3 Quality Analysis Using Standard Deviation of Peak Frequencies

In this quality analysis, the demodulated output is run through a peak detector. The standard deviation of each of the peak frequencies is calculated and the results are averaged over the range of the sweep. Again, in trying to identify noise verses audio or voice, one would expect a properly-demodulated signal to have a higher standard deviation. The peaks are distinct and their frequency distance from each other is relatively higher than that of a flat noisy curve. The results of this experiment are presented in Figure 4.13 and Figure 4.14.

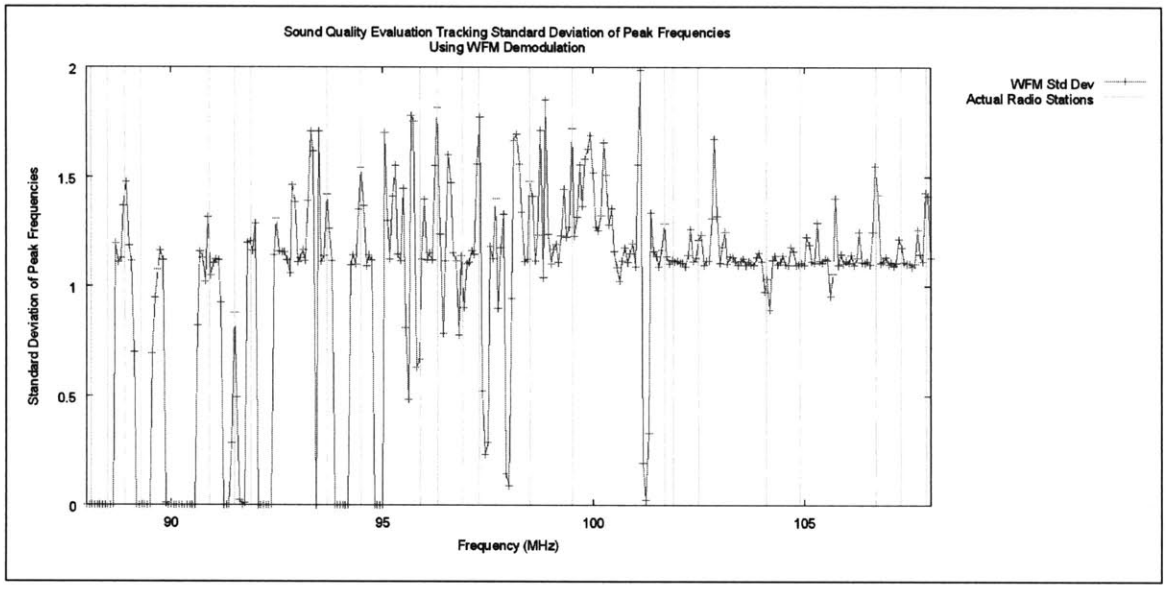

Figure 4.13 Brute-Force Demodulation Quality Using WFM and Peak Frequency Average Standard Deviation

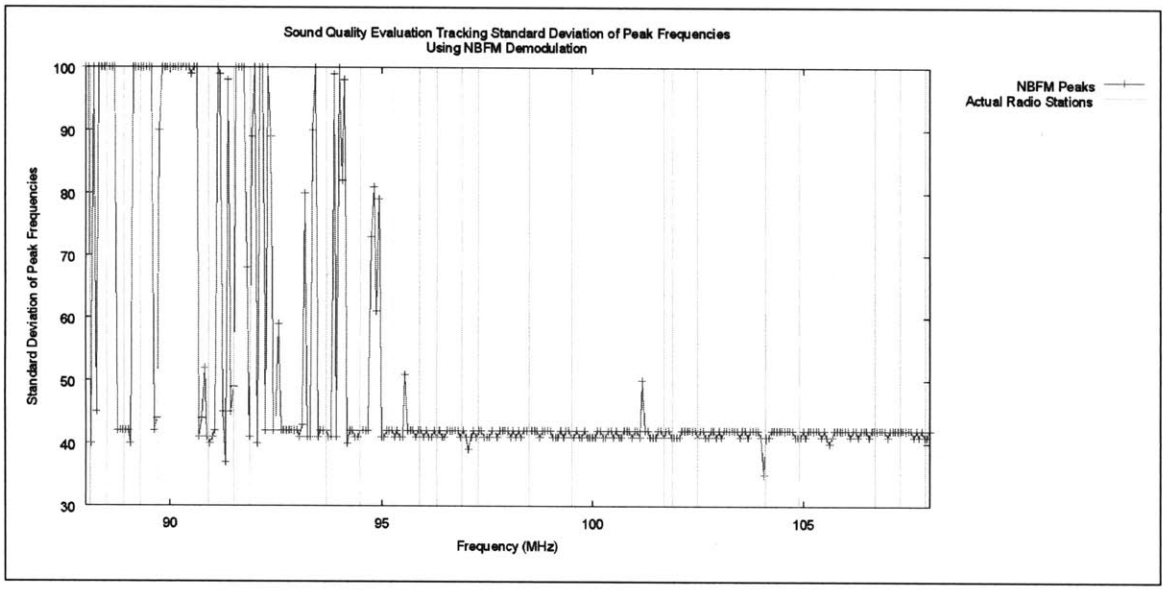

Figure 4.14 Brute-Force Demodulation Quality Using NBFM and Peak Frequency Average Standard Deviation

Although I expected this quality estimate to perform the best among the three used in these brute-force detection experiments, it seems to be the least reliable of the three techniques. It is generally clear that the standard deviation for the NBFM case is much higher than that for the WFM case, but this is the opposite of what was expected.

There are no doubt more elaborate schemes for assessing audio quality than the basic ones applied here which may well do a much better job at identifying a properly-tuned station.

4.6 Detecting Listeners

Rather than build a sweep application to try to generally detect listeners, I have built an application which can spot listeners to a specific, specified radio station. In Figure 4.15, the top graph indicates that the program is set to spot listeners in the immediate vicinity of 90.9 MHz. The frequency spike of the local oscillator is clearly

visible at about 175 kHz above 101.60 MHz, or 101.775 MHz. The other spike at around 100 kHz below 101.60 MHz is DC noise generated by the USRP Hardware.

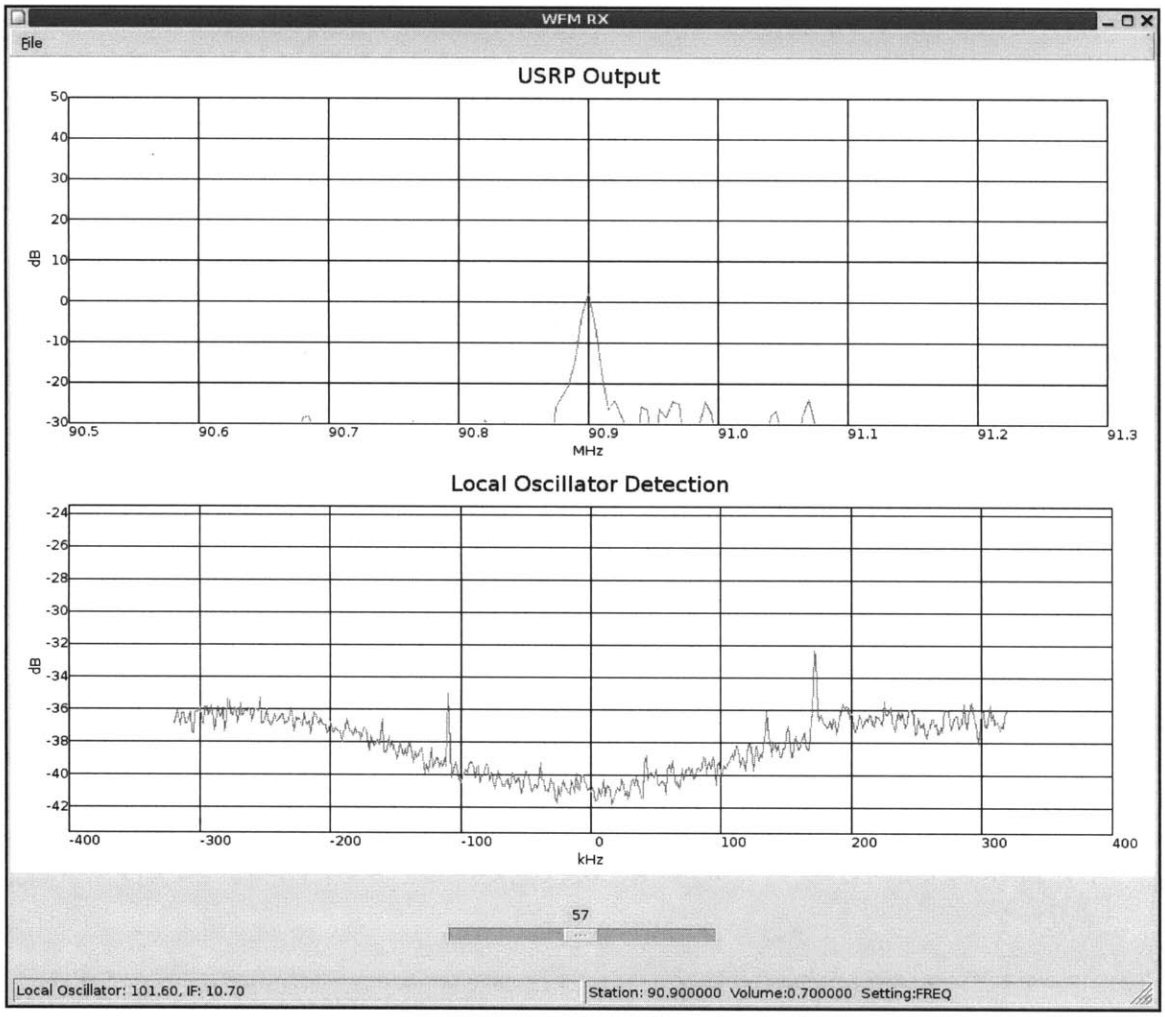

Figure 4.15 Local Oscillator Detection to Detect a Listener of a Specific, Known FM Radio Station

Conclusion

In this thesis, we have constructed software radio algorithms which might be applied in cognitive radio systems to make better used of rf spectrum space. First, we reviewed radio and regulatory history to try to understand the legacy allocations viewed as prime candidates for spectrum reuse. And, we reviewed the work that has been done so far in this area. Next, the work done for this thesis is placed in the context of how it relates to cognitive radio technology. In the Method section, we review specific software-defined-radio modules that were designed and constructed to build better cognitive radio systems, that can begin to understand the nature of existing spectrum occupants and plan reuse accordingly. Finally, we evaluated the output of each software module.

Future work in this area that has not been done for this thesis would be to expand upon the general themes and areas of detection as outlined in Chapter 2. The modules implemented here merely represent a simple prototype of what such sensing can be

included in cognitive radio systems. It is a powerful and important idea to attempt to build real systems using advanced sensing and analysis capabilities available in the RF domain.

The code-base for GNU Radio software and for this project is large. The main blocks of GNU Radio code have approximately 150,000 lines of non-commented C/C++ code and 17,000 lines of Python code. For this thesis, most of the work was done using Python. There were about 9,000 lines of Python code, 4,000 lines of C/C++, and 300 lines of Perl (the Perl was used for creating the NBFM Graph). All graphs were created using gnuplot software.

Worth noting are the design challenges associated with working in the emerging field of software radio engineering. The fact that we design and plan software-defined radio modules in pipeline blocks is a departure from decades of radio hardware design standards and requires that one understand existing designs but also depart from them as and when better solutions can be implemented in software. Many of the "tricks" and architecture principals that have been learned and refined since the dawn of radio may not necessarily be appropriate for software radio engineering. This is an emerging discipline and is warrants further study.

Finally, this thesis was intended to illustrate how rich an opportunity lies before us. The title, "A Day in the Life of the RF Spectrum" was chosen to particularly suggest an entirely new way of thinking about this domain as a whole, without divisions as has traditionally been done. While we hardly begin to study this vast resource as a whole in this thesis, we hope to encourage future work to think in equally open, broad terms as technology now permits.

Bibliography

[1] Weightman, Gavin. Signor Marconi's Magic Box. Cambridge: Da Capo Press, 2003. 86.

[2] *Ibid.* 142.

[3] http://www.oldradio.com/archives/jurassic/marconi.htm

[4] http://www.newscotland1398.net/nfld1901/marconi-nfld.html

[5] Sobel, Robert. RCA. New York: Stein and Day, 1986. 20.

[6] *Ibid.* 25-30.

[7] *Ibid.* 39.

[8] *Ibid.* 48.

[9] *Ibid.* 60-70.

[10] E. Krasnow, L. Longley, and H.Terry. "The Politics of Broadcast Regulation." 3d ed. 1982. Rpt. In Regulation of the Electronic Mass Media: Law and Policy for Radio, Television, Cable and the New Video Technologies. Ed. Douglas H. Ginsburg et al. St. Paul: West Publishing. 1991. 20-21.

[11] Sobel, 119-120.

[12] *Ibid.* 129-131.

[13] *Ibid.* 159-162.

[14] *Ibid.* 163-167.

[15] National Telecommunications and Information Administration. United Status Frequency Allocations: The Radio Spectrum. October, 2003. http://www.ntia.doc.gov/osmhome/allochrt.html

[16] Klemperer, Paul. "The Wrong Culprit for Telecom Trouble." Financial Times 26 November, 2002: 21.

[17] McHenry, Mark. Frequency Agile Spectrum Access Technologies. Shared Spectrum Company. Presented at FCC Workshop on Cognitive Radios. May 19, 2003.

[18] Marshall, Preston. XG: Next Generation Communications. DARPA. Presented at FFC Workshop on Cognitive Radios. May 19, 2003.

[19] Mitola, Joseph. Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio. Dissertation, Doctor of Technology, Royal Institute of Technology, Kista, Sweden. 2 May, 2000. 1

[20] *Ibid.* 59-50, 289.

[21] FCC. ET Docket No. 03-108. March 11, 2005.

[22] FCC. ET Docket No. 04-186 and ET Docket No. 02-380. May 25, 2004.

[23] FCC. Spectrum Policy and Technology: Spectrum Access and the Promise of Cognitive Radio Technology. Presented at the FCC Workshop on Cognitive Radios. May 19, 2003.

[24] Mitola, 59.

[25] http://www.gnu.org/software/gnuradio/

[26] http://www.ettus.com

[27] http://www.fftw.org/

[28] http://sourceforge.net/projects/pygarmin

[29] http://pyopengl.sourceforge.net/

[30] International Telecommunication Union. Recommendation ITU-R BT.470-6, Conventional Television Systems. 1998.

[31] ElBoghdady, Dina. "Advertisers Tune In to New Radio Gauge." The Washington Post 25 October, 2004. E01.

Appendix A: Code Modules

The code modules and project description constructed for this thesis are available on the Viral Communications Web Site, http://viral.media.mit.edu