

Characterization of Spectrum Activities in the U.S. Public Safety Band for Opportunistic Spectrum Access

Steven D. Jones, Eric Jung, Xin Liu, Naim Merheb, and I-Jeng Wang

Abstract— Dual site spectrum measurements have been made in the public safety band in Howard County, Maryland, USA. The public safety band is of interest because of its obvious importance and the increasing need to determine whether improved spectral utilization would accommodate increased usage for public safety. Two receiver suites were synchronized and used to measure concurrent spectra at separations between a few meters up to a few kilometers. These measurements prove useful in examining issues of spectrum sensing for dynamic spectrum access, including receiver sensitivity, primary user signal detection, adjacent channel interference, and policy performance with local and cooperative sensing. We analyze the collected data to characterize the usage of this public safety network and provide insights into how white space can be identified and utilized. We propose a class of opportunistic access strategies and demonstrate how collected data can be used to evaluate the performance of such schemes.

Index Terms—dynamic spectrum access, public safety band, spectrum measurement, spectrum policy, and spectrum data analysis

I. INTRODUCTION

Wireless communication faces constant challenges from the ever-increasing demand for wireless services requiring additional spectrum. It has been widely recognized that there exists a large amount of white space (unused or lightly occupied spectrum) in time, space, and frequency. It is often the absence of dynamic channel access instead of the true spectral scarcity that limits the growth of wireless communications systems [1,2,3]. With this realization, federal government agencies, as well as industrial interest groups, have taken important initiatives towards flexible and dynamic spectrum policies, including regulation recommendations, secondary market spectrum leasing rulings, and technical model proposals. At the same time, advanced semiconductor and Radio Frequency (RF) technologies have produced devices that are more intelligent and less expensive with stronger sensing and signal processing capability. Driven by necessity and enabled by new technological advances, now is the perfect time to develop cognitive radio networks for dynamic spectrum utilization.

To enable dynamic spectrum utilization, it is critical to understand the spectrum usage of primary users and identify the existence/absence of primary users and policies for secondary access. These needs motivate our measurement and analysis on the public-safety band in this paper. We address this problem by examining spectrum data measured in the public safety band at two locations concurrently. The objectives of the measurement study are as follows: First, it provides a real sample of spatially diverse spectrum activity data, which can be used to quantify primary users' spectrum utilization factor and access pattern, and to evaluate current proposed schemes on channel probing and user detection. Many practical issues arise here. Second, the measurement data enables us and other researchers to study spatial correlation and temporal correlation between synchronized sensing units and across time. Potentially, a good model for channel activities can be developed and used to generate synthetic data for further testing and study.

We choose the public-safety band because of its importance and to determine whether improved spectral utilization in the band would accommodate increased usage for public safety. The primary system in this public safety band is a trunking system with a county-wide set of repeaters. The public-safety band in general has a high peak usage (at the time of emergency) and low average usage, and thus has been considered by the Federal Communication Committee (FCC) as a potential candidate for interruptible secondary usage with stringent primary user protection. Such characteristics are of interests to dynamic spectrum access studies. We also note that the secondary users in the public-safety band can be other public-safety users (with potentially lower priority) instead of unlicensed devices. We acknowledge that different spectrum bands at different locations and different time can have different characteristics. For example, the features of a TV band may be vastly diverse from a public-safety band. Therefore, our measurement study is one realistic spectrum utilization trace to be used to study potential dynamic spectrum utilization.

In this paper, we present a measurement-based study of the public safety band at Howard County, Maryland. Two synchronized spectrum analyzers are set up to conduct simultaneous measurements at different locations. We conduct

six sets of measurements at six different pairs of locations in two days. We use the two measurement sets to study the spatial correlation between sensing units separated by distances of 5 meters to 6.5 km. In each measurement set, we collect spectrum activity data for a 20MHz band for 30 seconds for the forward and the reverse channels.

We configured two nearly-identical data collection suites. Each of the suites includes an omni-directional antenna, filters, amplifiers, downconverters, analog-to-digital converters (ADCs), storage devices, and timing devices. The ADCs sample at 64 MSamples/second periodically according to a pre-arranged schedule. This schedule is designed to ensure synchronized collection at the two sites. Global Positioning System (GPS) clocks are used to maintain synchronization between the two suites. The clocks provide a trigger that is used to begin a sample data collection independently but simultaneously at the two sites.

Based on the measurement data, we conduct a set of analyses. We identify busy and idle channels, as well as channels subject to adjacent-channel interference. We analyze the impact of the detection threshold used to identify channel activities. We also study the correlation between signals at two locations and the impact of device sensitivity.

In addition, we propose an opportunistic access strategy based on a sensing period and a keep-off period. Using the measurement data, we study the impact on secondary spectrum utilization and interference to primary users as a function of the parameters of the proposed policy. We also compare the result from the two measurement sets.

Dynamic spectrum utilization has attracted a lot of interest. Measurement studies have been conducted by various groups at different locations and time, e.g., [1], [2]. These studies indicate the (wide) existence of white space while limited analysis has been conducted based on such measurement data. Sensing-based channel probing and user detection has attracted a lot of attention due to its simplicity and low requirement for system infrastructure support, e.g., [4]–[6], [8]. Most of these works are based on theoretical analysis and channel models. A consensus from these studies is that cooperative sensing is critical to the success of sensing-based spectrum access due to the existence of severe fading and shadowing. While a large amount of theoretical analysis and protocols have been proposed for sensing-based approach, little is known on the applicability of the schemes in reality. Our measurement can be used to validate the theoretical analysis and schemes proposed in this research area. The most related work is the measurement study in the cellular band conducted by Shah et. al. In [9], the authors present a sensing architecture for wideband measurement and distributed measurement method. The authors collect measurement data in a GSM band and apply a decision fusion scheme to quantify the spectrum usage. In comparison, our measurement study focuses on the public-safety band with different characteristics. Furthermore, we focus on analyzing the

realistic implications of the measurement data on different schemes and study correlation between measurement sets. There are also a lot of research efforts on centralized and decentralized spectrum sharing/auction schemes, and medium access control protocols, e.g. [7], [10]. Most of these works assume given channel activities and rely on some kind of detection schemes. Therefore, the measurement of real network activities provides an important step toward a realistic understanding of dynamic spectrum utilization.

The paper is organized as follows. A brief description of the US public safety network under study is presented in Section 2. We present the sensing system and data collection methodology in Section 3. By analyzing the spectrum data, we have gained some interesting insights, as discussed in Section 4. In Section 5, we present the analysis of an opportunistic access strategy. We conclude the paper in Section 6.

II. THE US PUBLIC SAFETY NETWORKS

A representative legacy public safety network is the land mobile radio systems currently in operation. Howard County, Maryland operates a trunked land-mobile radio system in the 800-MHz band normally assigned to public safety networks. The system's infrastructure consists of nine repeater sites networked by microwave and terrestrial links back to a central controller. The system supports roughly 2048 talk groups with twenty-three virtual channels. A virtual channel is actually a pair of frequencies separated by about 45 MHz, where one frequency in the pair is used for the forward channel (repeaters to subscribers) and the other for the reverse channel (subscribers to repeaters). Traffic from one subscriber to others in the same talk group always passes over the air twice with the repeater acting as a relay, using two different frequencies. The forward channel center frequencies are in the range 856.2375 to 868.0625 MHz and the reverse channel center frequencies are in the range 811.2375 to 823.0375 MHz. See [11] for a broader description of the public safety band. One of the virtual channels is assigned as a control channel on a daily basis. Traffic consists entirely of FM voice with a bandwidth of less than 20 kHz, except for the control channel that carries FM data in a similar bandwidth.

Subscriber units are either vehicle mounted or personal and each is assigned to one or more talk groups. A subscriber selects a talk group then keys the unit. The unit signals via the daily control reverse channel its intentions to the central controller, which assigns a virtual channel to the talk group. This assignment is broadcast to all members of the talk group via the daily control forward channel. Once the assignment takes place, the subscriber units tune to the newly assigned virtual channel and the requester is allowed to contribute on the corresponding reverse channel. All repeaters are simultaneously receiving all channels at all times but the central controller votes on which one is instantaneously receiving the best signal from the current contributor. It is this best reception that is sent to other subscribers in the talk group via the forward channel in the assignment.

The system does not use knowledge of the location of subscribers such that traffic is transmitted via the forward channel simultaneously at all repeaters, contrary to a cellular telephone system. As a result, transmissions on the forward channel can be detected throughout the entire county, while reverse channel transmissions are confined to the locality of the emitter and the nearest repeater. This presents an interesting opportunity for dynamic spectrum reuse since there will be far more occupancy in the spectrum of the forward channels than in that of the reverse channels at any given location.

III. SENSING SYSTEMS AND DATA COLLECTION METHODOLOGY

We sample the virtual channels with two synchronized receiver suites. Each suite is composed of an antenna, pre-selection filter, low noise amplifier, downconverter, A/D converter, and a GPS clock. The two suites are synchronized in time and triggered by a common script so that spectral measurements are made concurrently. In some cases, one suite is tuned to the spectrum of the forward channels, while the other is tuned to the spectrum of the reverse channels to investigate frequency dependencies. In other cases, both suites are tuned to the same spectrum to investigate space dependencies. The spectral data are sampled with 14 bits at 64 MS/s. The samples are collected in snapshots of 16384 each, covering ~ 20 MHz every 100 ms. This snapshot size allows for DFT bins of ~ 3.9 kHz, which is enough resolution to discriminate between adjacent forward or reverse channels as their centers are never closer than 25 kHz.

Spectral data were collected over two days at separation distances between the two suites of 5, 200, 500, 2300, 3800, and 6500-meters. In general, one suite remained fixed at a selected site, while the other suite was positioned at the distant locations. In the sequel we refer to the sensor at the fixed site as site 1, and the sensor at the mobile site as site 2.

IV. ANALYSIS AND INSIGHTS

In this section, we investigate the collected measurement data to identify properties that can aide us in characterizing the usage of this public safety network. The objective of the analysis is to provide further insight into how white space can be identified and utilized. In general, the data sets allow us to understand how these channels are utilized at the time of measurement, and the effects of activity in the spectrum surrounding these channels. We also investigate the effectiveness of a threshold-based detection method based on the data and spatial correlation between the two measurement sets. In this analysis, we primarily consider the forward channel band measurements because of the very localized nature of the reverse channel activity.

A. First Impressions

At first glance, the spectral power maps of the channels provide a fairly clear indication of the channel occupancy.

Figure 4.1(a) shows one such map for the entire 20 MHz spectrum band (from 850MHz to 870MHz) at site 1 for the second measurement set, using a color gradient to convey signal power observed in the spectrum at each time. High power usage is easily determined by the colorscale, and the frequency of strong transmissions can be determined in the graphs, especially because in this case they occur seldom enough to draw strong contrast to the lower power or completely unused spectrum.

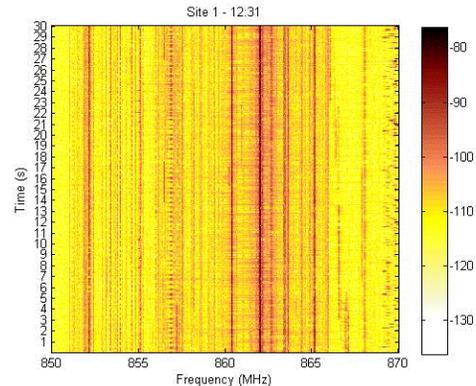


Figure 4.1(a): Spectrum density map for 850-870 MHz band

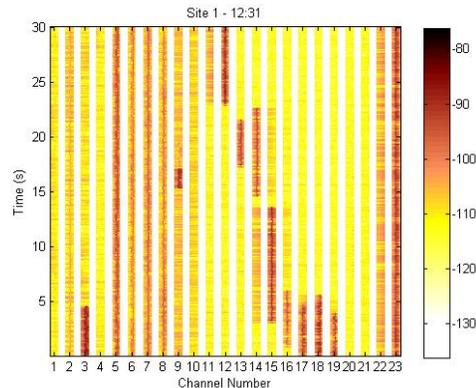


Figure 4.1(b): Spectrum density map for the 23 virtual channels of the public safety band

Note that the 23 virtual channels for public safety users are unevenly spaced within the 20MHz band, and that each channel only has a bandwidth of 25 KHz. Further, the actual center frequencies for the channels are actually within the 856-869 MHz band. Parsing out the frequency bands used for the public safety channels and using this same color scale, we can then observe which channels show strong activity in each measurement set.

Figure 4.1(b) shows the spectral power maps for the public safety bands in the measurement set from figure 4.1(a). The white strips in between are used to separate channels and are used throughout this section. We can see clearly that channels 5 and 23 are occupied for the complete duration of the measurement set, and some strong activity is seen for limited

durations in various other channels. There is also a great deal of weaker activity in some other channels, which through methods explained later, can be determined to be either real activity or interference. We also observe that Channel 23 is consistently strong and active in all measurement sets. Our conjecture is that Channel 23 is the control channel during the measurement duration.

B. Quantitative Analysis and a Threshold-based Method

In order to measure channel occupation quantitatively, we define a channel as being active at a given time if the power level of the channel at that time is greater than a threshold (τ_e). The threshold is set above the noise floor measured for that particular set. The noise floor used for each set is essentially an average of all of the power data points over the set. As Figure 4.1(a) shows, a great majority of the 20 MHz spectrum is inactive, so a reasonable noise floor number can be obtained using this method. Noise floor numbers in general will not be definite for a given spectrum because of the randomness of environmental variables.

Since the resolution of the data at each time interval is 3.9 KHz, and the bandwidth of a channel is ~25 KHz, there are six to seven data points for each channel at each interval. Thus, using the seven data points centered at the center frequency of each channel, we can consider that channel active at that particular time if any of these seven bins is greater than the threshold. Figure 4.2 shows how this definition is applied to evaluate the activity in the channels.

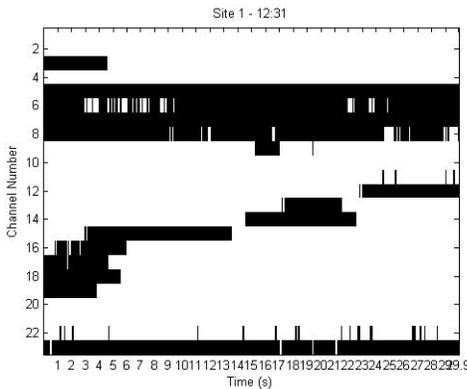


Figure 4.2 Channel occupancy given a 10 dB threshold at site 1 at time 12:31 PM.

In the figure, the threshold is set to 10 dB above the noise floor, using a black tick to represent activity in a particular channel at a given time, for the measurement set shown in Figure 4.1(b). Alternatively, we can also evaluate the sum of the energy in all bins when applying the thresholds.

In such a threshold-based method, choosing an appropriate threshold is extremely important in evaluating the availability of space in the set of channels. It is important that we find a range of threshold values that enable us to distinguish between obviously active, less active, and completely idle channels. To gain insight into how this threshold might be found, we use

plots of how long each channel is detected to be occupied over a range of thresholds.

C. Active and Idle Channels

Comparing the channel maps from Figure 4.1 with the percentage occupation graphs of Figures 4.3 (a)-(c) lends insight into the nature of active channels. We can see heuristically from the power map in 4.1(b) that channels 5 and 23 are active over the entire duration of the data set, with some strong limited duration activity occurring in channels 13-19. Figure 4.4(a) shows us curves that correspond to some of the channels that show strong activity. The upper curves represent channels 5 and 23, and the lower ones represent channels 13, 15, and 17.

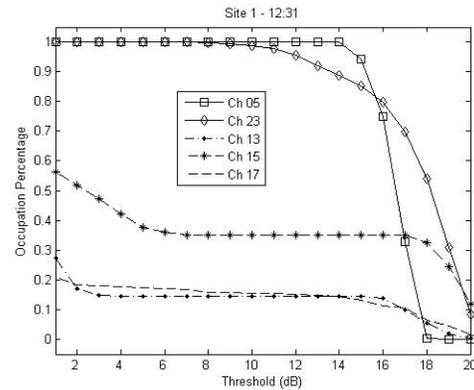


Figure 4.3(a) Channels with strong activity

We first observe that these channels appear to stay level consistently over a wide span of mid-range threshold values. Channels 5 and 23 are clearly active over the entire measurement period, while the lower level curves indicate strong activity in their respective channels for only a fraction of the measurement duration. As the threshold climbs above 15 dB, the percentages begin to drop sharply, as the threshold level becomes too high for even the strongest signals to cross. **We consider the consistency over a wide threshold range and the severity of the drop as good characteristics in determining whether a channel is active or not.**

Using these observations, we can also classify some of the channels that exhibit less intense activity as occupied channels. In particular, channels 6-8, which appear to have less intense power levels across the entire data set duration, display in Figure 4.3(b) a fairly consistent occupation level over the same range. We found that these same trends are observed in all of the data sets. We may wish to apply a frequency-dependent threshold to address the disparate nature of the channel activity.

But what are we to make of the other channels? Figure 4.3(c) shows the occupation percentage curves for some of these remaining channels. We see that some of these curves, e.g., channel 20, dip directly to zero outside of the low threshold range (1-6dB), and some, e.g., channels 9 and 14, drop sharply from much higher percentages.

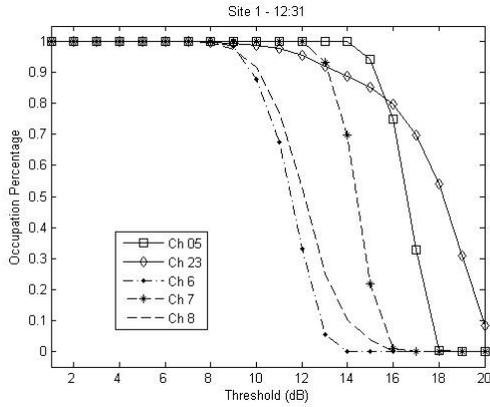


Figure 4.3(b) Channels with good activity

Comparing with their corresponding power maps in Figure 4.4(b) shows that of these remaining channels, some appear to have no activity whatsoever, and others appear to have much lower power activity than the clearly active channels. The latter requires further investigation regarding the causes of their activity, whether it be intentionally low power transmission or simply interference from neighboring sources. We can then form a stronger basis for their classification as empty or occupied.

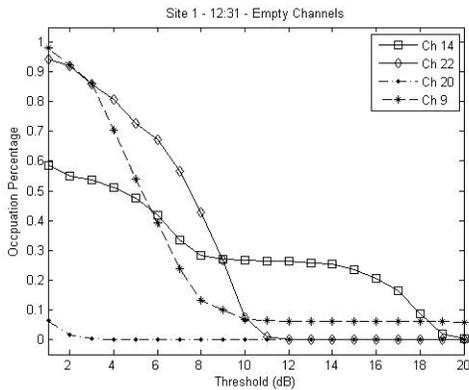


Figure 4.3(c) Channels with weak activity

In summary, we observe that the consistency of channel occupation over a range of thresholds and a sudden drop as good characteristics to determine an appropriate threshold. The range of the threshold where occupancy is consistent depends on the activity strength. The stronger the activity (i.e., high power), the longer the consistent range. The sudden drop is also a good characteristic to distinguish activity from noise.

D. Adjacent-channel Interference

Investigating the spectrum surrounding these low power channels reveals that their activity is caused by interference from adjacent bands (including both public-safety channels and non-public-safety channels). The most severe cases of this interference can be seen in channels 14 and 22, which lie adjacent to channels 15 and 23, respectively. Figure 4.5 presents the power maps for channels 14 together with 15, 22, and 23 with their adjacent 200 KHz bands. The adjacent 200

KHz bands in these figures are not composed of channels in the public safety spectrum. In Figure 4.5(b) we can see that the activity in channel 22 results from interference from the high power transmission in channel 23, and we see a closely mimicked pattern in the spectrum band immediately greater than channel 23 as well.

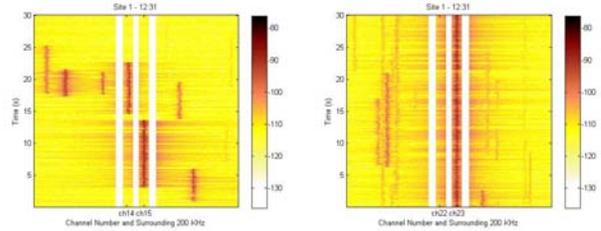


Figure 4.4 (a) Adjacent channel interference between channels 14 and 15 and (b) Adjacent channel interference between channels 22 and 23

This is also the case with channel 14 in Figure 4.5(a); although we see strong activity in channel 14 between ~14-23 seconds due to a transmission, the much lower power activity in the channel between ~4-13 seconds correlates exactly with the strong transmission in channel 15. We see that channel 14 sees a gradual fall in occupation percentage similar to that of channel 22, until it reaches a steady level due to its strong activity at a threshold level of ~8 dB. We also observe similar behaviors in the public safety channels with interference caused by adjacent non-public safety channels. For instance, the same pattern is seen in channel 9, which lies near a strong transmission in the surrounding non-public-safety spectrum.

The result of this investigation highlights a tradeoff regarding how the threshold should be chosen. The occupation percentage maps show that the 1-5 and 15-20 dB ranges can be ruled out as too conservative and too aggressive to accurately reflect the actual occupation of the channels. A desirable threshold should be selected in the mid-range and depend on how we wish to interpret interference. We can see that a higher threshold within that range will negate adjacent channel interference as intentional usage of the compromised channels. On the other hand, a lower threshold will assess interference in the channels as activity, which, although not an accurate reflection of access by the primary users, may be a more accurate indication of the actual usability of these channels by a cognitive radio device. This highlights an important challenge for low-cost sensors. The sensors must have high dynamic range to avoid saturating when strong primary signals are detected, and tunable filters to isolate sensed energy in specific channels.

E. Correlation between the Two Measurement Sites

Figure 4.5 shows the spectral power maps for the 23 public safety channels for the first data set at both of the measurement sites. During this measurement period, the vans were only 5 meters apart and were both stationary. There are a few observations we can obtain from these maps. The first is

that the two sites observe very similar activities across the measured spectrum. On the other hand, site 2 appears to detect stronger activity compared to that of site 1. The latter observation may be attributed to differences in the sensitivity of the two measurement suites. Although site 2 does detect stronger signals, in Figure 4.5 the similarity in activity in each of the channels is present. This is to be expected, of course, given the close proximity of the two measurement sites.

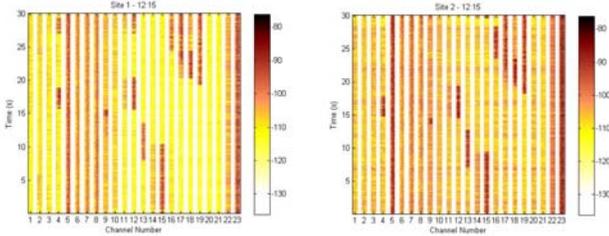


Figure 4.5 (a), Spectrum density at site 1, and (b) Spectrum density at site 2

Applying a simple threshold to these two sets shows that the data yields fairly different levels of correlation for different channels. Figure 4.6 shows an occupation comparison graph between these two sites at a threshold of 10dB above their *respective* noise floors. The white and black marks show agreement in the detection of empty and occupied space respectively between the two sites, while the grey ticks show a mismatch between the sites at a given time in a given channel. We can see that, in particular, channels 6 and 8 appear to differ greatly because site 2 detects a large absence of activity while site 1 detects that the channel is occupied for the entire measurement duration.

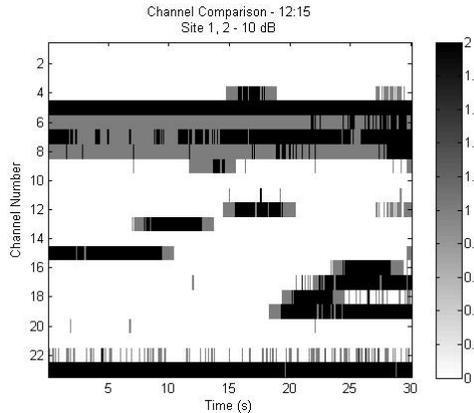


Figure 4.6: Comparison of the activities at the two stations at a threshold of 10 dB [scale is the number of sites detecting]

This observation presents an interesting practical issue. *These two stations have set the same threshold levels, but have yielded very different results in detection.* Given the fact that two experimental spectrum sensors can have different measurements and should have different values of threshold, it is very conceivable that (low-cost) spectrum sensing devices can have different sensitivities.

Therefore, it is of a practical interest to integrate the

measurements of different sensing devices. Possible solutions depend on application scenarios. If the sensors have access to high capacity links, it may be desirable to report all measurement data to a central controller and use a soft-detection algorithm which considers potential device discrepancy. Another common case is one in which the sensors are communicating with each other through a limited control channel. In this case, an interesting problem would be to understand how to adapt to device sensitivity and choose an appropriate hard-decision scheme (potentially different on different devices).

In the future, we would like to study the degree of correlation between measurement sets at varying distances. In addition, there are many possible causes of differences in measurement stations, such as 1) fading environment; 2) sensitivity/calibration differences in devices; 3) distance to a repeater, and 4) the calculation of the noise floor. We consider the cause of the difference in this data set as sensitivity/calibration because site 1 shows a consistently stronger signal than site 2 in all measurements at different locations. However, we would like to better understand the issue and design further experiments to distinguish some of the aforementioned causes.

V. EVALUATION OF OPPORTUNISTIC ACCESS STRATEGIES

With the available public safety band spectral data described above we now examine access strategies based upon policy. This analysis represents one example of how these strategies may perform when put into practice. Our data, although representative of the public safety band, does not give us a ground-truth indication of primary user signals, is limited to a 30-second horizon, and was collected during a nominal traffic time (i.e., not during a significant public safety incident). A more comprehensive analysis will be valuable.

Consider a spectrum sensing suite integrated with a radio capable of dynamic spectrum access. Further consider that there are two such radio/sensor suites (or cognitive radios) which are separated by some distance and have the objective of communicating with one another in the public safety band.

With the ability to measure spectrum at each site, this pair of radios may employ a variety of policies in determining how to use the available spectrum while minimizing the interference to primary users. In this section we consider the impact of local sensing in which each radio measures the local spectrum and decides on the availability of whitespace, and coordinated sensing in which the pair exchange spectrum data measured by each, and make decisions based upon the combined spectrum data. In each case, we examine the potential secondary throughput and resulting interference on primary users at each site (see Figure 5.1).

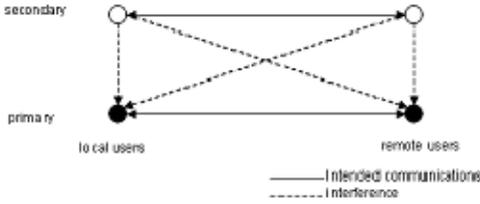


Figure 5.1 Assumed topology of the primary and secondary user(s)

Secondary channel utilization is calculated the percentage of all time and frequency slots that can be occupied by a secondary signal in accordance with a defined policy. We also define a secondary utilization normalized by the available whitespace, indicating the success of the secondary user in accessing the whitespace, exclusive of interference. Interference is measured as the fraction of all time slots in which the secondary user transmits while the primary user signal co-exists in the same time slots. Normalized interference is defined as the fraction of primary user timeslots in which interference occurs. This measure is important because as primary user access decreases, the possibility of interference also decreases. Referencing the interference to the primary throughput is useful in this context.

Given that the ground truth of primary communications cannot be completely recovered from the spectrum activity in measurements, we assume that the sensing is perfect when evaluating the secondary access strategies. That is, primary user signals are presumed to be present whenever the spectrum sensing suite detected energy above a threshold in the channels of interest. This is obviously an approximation to the reality. Methods for selecting a threshold have been described in Section IVB.

A. Opportunistic Spectrum Access Strategies

We define a class of strategies for the secondary users that can be characterized by three parameters: **the energy threshold**, **the spectrum sensing period**, and **the keep-off time**. We first define these parameters and briefly discuss qualitatively their expected impacts on performance:

The energy threshold (τ_e) is applied to the spectrum measurements to differentiate whitespace from primary user signals as described in Section III. The energy threshold used at each sensing site must be chosen on the basis of probability of false alarm and probability of missed detection. A low threshold will result in best signal detection, but at the cost of false alarms due to noise. These false alarms will decrease secondary utilization of the spectrum, as the secondary user will inhibit transmitting when the false alarms are detected. Conversely a high threshold will avoid false alarms, but may result in missed detections of true primary user signals. When this occurs the secondary radio may transmit during primary user transmissions, causing interference. When the primary signal is presented at the sensor with low power due to distance, fading, or low transmit power, the sensor is best

served by a relatively low threshold, particularly if the intended primary receiver is closer (or in a lower loss path) to the secondary radio as in a hidden node configuration. Figure 5.1 shows this possibility in that the local secondary user must detect the remote primary user so as to avoid interfering with the intended communications from the remote primary user to the local primary user. Based on the findings in Section IVB we have chosen a threshold of 5 dB above the mean noise floor of each of our receivers to demonstrate the policy concepts herein.

The spectrum sensing period (T_s) defines the rate for periodic sensing that takes place concurrently on each channel of interest. This allows the secondary sensor to check for the presence of a primary user signal on a frequent regular basis. In the collection of public safety spectrum data in the field we collected 256 μsec of spectrum data every 100 msec. Thus for (100 msec – 256 μsec) our sensor is idle. This was done to save data storage space. In evaluating various sensing period policies (T_s), we select the sensed data at the corresponding periods from our data, with the minimal period being 100 msec. The intervening spectrum data is assumed to be unavailable to our sensors and is ignored. We also ignore the (100 msec – 256 μsec) period after the spectrum is sensed. In essence, we assume that the sensing requires 100 msec. In general, one would expect frequent spectrum measurements to track primary user activity closely and enable the secondary user to rapidly detect a returning primary user and thus abandon the channel quickly, minimizing interference. The secondary user, however, does not transmit while sensing so as not to interfere with the sensing.

The keep-off time (T_k) is defined as the minimum time duration over which no primary activity is detected before the secondary users access a channel. The secondary user may not access a channel until it has been observed to be idle for this period. The keep-off time is designed to increase the probability of detecting primary activities. When primary user activities on a channel are sporadic, it may be valuable for the secondary users to verify that the primary user has vacated a channel for a pro-longed period before the secondary user decides to use it. Under the constraint on sensing rates, this can be accomplished by repeated sensing. Additionally, when the primary user signal is weak at the secondary user sensor, a missed detection can occur. If this is the case, repeated sensing can reduce the likelihood of missed detection, thus mitigating interference. Since the sensing is assumed to be periodic (for now), the keep-off time can be specified equivalently by selecting the appropriate **keep-off counter** N_k . That is, a secondary access is attempted only after N_k sensings that led to no detection of primary communications. If a detection is made, the counter is reset to zero. In this case, the corresponding keep-off time is $T_k = (N_k - 1) \cdot T_s$. In general, sensing need not be periodic, but should be coordinated among secondary users.

Figure 5.2 depicts the sensed activity on 23 public safety

band channels collected at a secondary radio site. A threshold of 5 dB above the mean noise floor was used to decide on the presence of primary signals. In some cases, noise or adjacent channel signals may have exceeded the threshold. In any case, this presents a basis for examination of the remaining policy parameters: T_s and N_k .

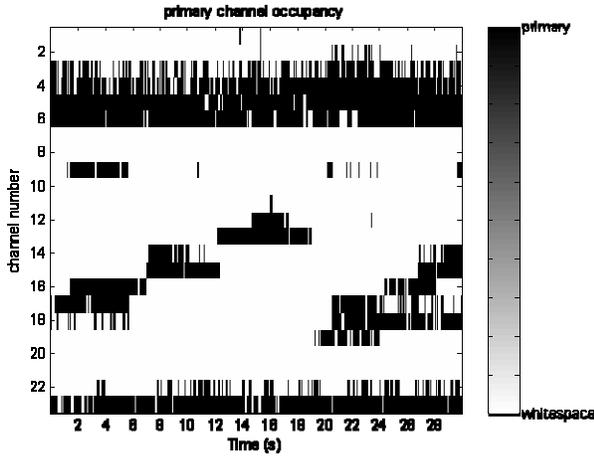


Figure 5.2 Detected primary activity at site 1

Figure 5.3 depicts the whitespace, primary accesses, potential secondary accesses, and resulting interference for a policy of $T_s = 1$ second, and $N_k = 3$ observations. Here the primary access is the same as in Figure 5.2.

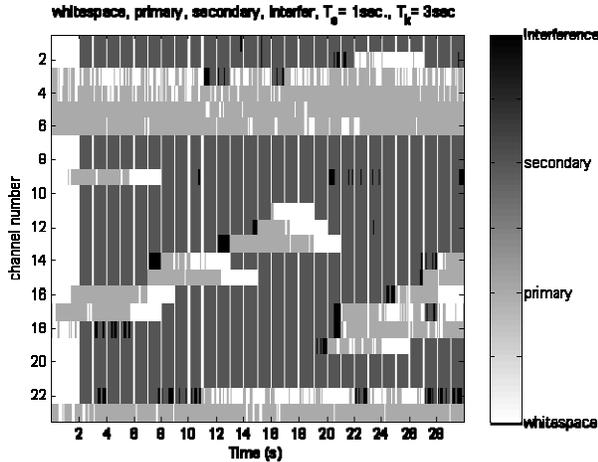


Figure 5.3 Characterization of spectrum accesses at site 1

Beginning at the time origin, the policy is applied to each channel independently, such that sensing occurs each one second (i.e., every 10th sensing period in the original data which was sensed at 100 msec intervals). Since the N_k requires 3 observations (at 0, 1, and 2 seconds), the secondary user will not start transmitting in any channel until the 3rd observation has occurred, even if the channel is entirely composed of whitespace. Thus the figure depicts a recurring sequence of sensing, incrementing of the keep-off timer for channels with no detection, and resetting of the timer when detections occur.

In Figure 5.3 it is apparent that primary transmissions are detected when the secondary access pauses to sense the spectrum since they are followed by whitespace for three sensing periods before the secondary transmission can again begin. Note that all viable secondary transmission opportunities are identified in Figure 5.3, to provide an upper bound on the secondary utilization. Secondary utilization of all timeslots, exclusive of interference events, is 51.2%. Normalizing, the secondary user is able to access 71.5% of all whitespace. Interference appears in Figure 5.3 when the primary signal reappears between secondary sensing periods and occurs in 1.7% of all timeslots, or in 6.1% of all primary timeslots. Since the primary signal is not detected until the next sensing period interference can occur until the secondary user abandons the channel. There are occasions when the primary signal is not detected at all due to its short temporal nature, and the secondary user would remain active on the channel. As noted, some of the primary activity in the figure may actually be due to false alarms resulting from noise, and hence interference would not occur.

B. Examining the impacts of T_s and T_k

We varied the choice of T_s over the range of 0.1 to 5 seconds to examine its impact on secondary utilization and resulting interference. Here we fixed T_k at 3 sensing periods. The result is plotted in Figure 5.4. When $T_s = 0.1$ second the secondary utilization is non-existent since every 100 msec period will be dedicated to sensing and the resolution of time in our measurements is 100 msec. As T_s increases, less time is used for sensing and secondary utilization increases until the T_k becomes large (note that T_k grows with T_s since it is $3 \cdot T_s$), incurring a penalty in the “slowness” of the secondary users to return to the channel after the primary departs. Eventually the time horizon of our data set at 30 seconds limits the secondary utilization as well. The best choice of T_s seems to be around 1 second for maximizing secondary utilization when $N_k = 3$.

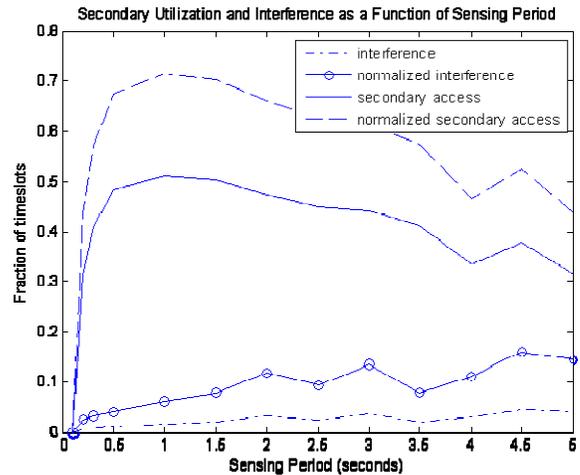


Figure 5.4 Impact of sensing period on secondary utilization and interference

As the sensing interval T_s grows, the idle time after a primary

detection increases. But interference also increases because the primary user is not detected as rapidly. The effect of T_k on secondary utilization is seen in Figure 5.5. Secondary utilization decreases with increasing T_k due to the increasing idle time upon detection of a primary user. Interference decreases slightly with increasing T_k due to the conservativeness in secondary transmission. Whitespace also persists longer when T_k is increased.

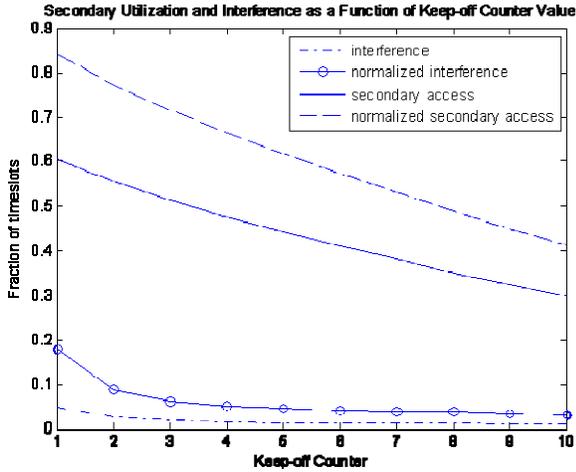


Figure 5.5 Impact of keep-off counter on secondary utilization and interference

Next we consider the impact of secondary transmissions from our site upon primary signals sensed at a site 3.8 km away. Sensors at the two sites detect differing spectra (compare Figures 5.2 and 5.6). Proximity to emitters and fading will result in differing signal levels being observed at each location. Similarity can be observed in the data from the two sites, however channel 6 is observed as occupied in one receiver while unoccupied in the other. The same is true of channel 7.

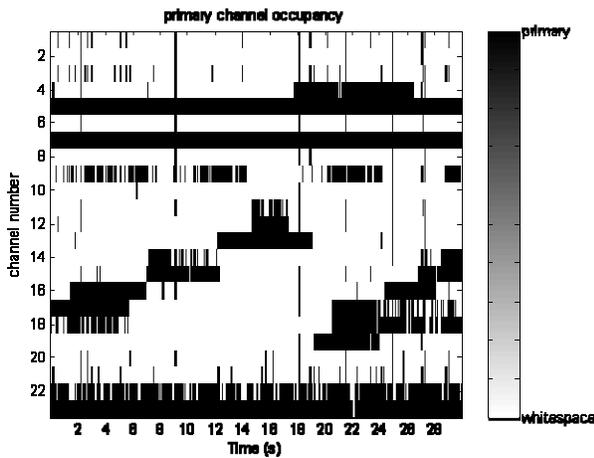


Figure 5.6 Primary detection with a 5 dB threshold at site 2

We assume that the secondary user at site 1 may interfere with primary users in the vicinity of either itself or site 2 if it elects to transmit on an occupied channel. When we consider local sensing only, that is, the secondary emitter selects a

channel and transmit only on the basis of spectral data obtained at its site, interference at the distant site may be found to be increased over the local interference. This is shown in Figure 5.7. This is due to primary transmissions at one site not being detected in some cases at the other site.

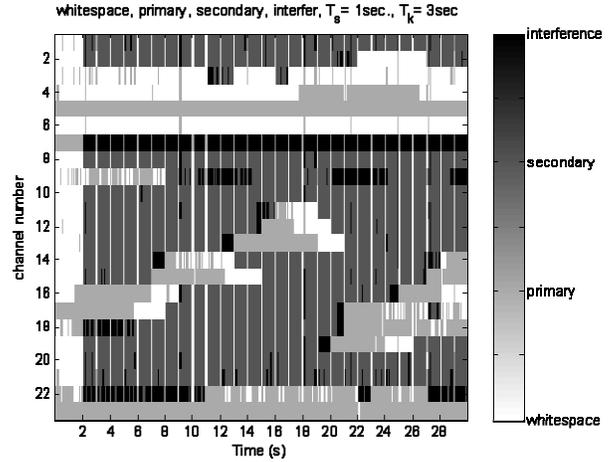


Figure 5.7 Local sensing only - site 1 secondary accesses super-imposed on site 2 primary accesses (8.7% interference), Site 2 secondary accesses super-imposed on site 1 primary accesses produce 7.2% interference (not shown).

Based on the observations of correlation between the spectrum sensed at each site in Section IVE, an obvious improvement would be expected with cooperative sensing, in which the two sensors share the detected primary signals. The decision on available whitespace is then made on the basis of the combined spectrum. Of course, this exchange of spectral data will require a common channel and will therefore reduce the secondary throughput available for user traffic. The result of cooperative sensing is shown in Figure 5.10. The reduced occurrence of interference can be observed as 2.1% in the cooperative case as opposed to 8.7% and 7.2% in the local sensing cases.

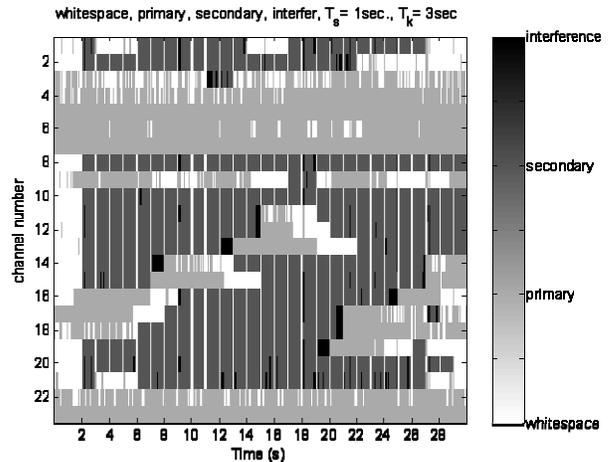


Figure 5.8 Result of cooperative sensing (2.1% interference)

C. On the impact of threshold

We next examine the influence of the energy threshold on secondary utilization and interference. We do not have a ground truth set for our spectral data collected in the field to indicate when the primary transmitter was active. We selected a threshold of 5 dB above the mean noise floor to represent an assumed ground truth in accordance with the findings of III B. Two other thresholds were then applied, 2 dB and 8 dB to the data of the previous section.

As expected, we have found that a lower threshold will make the secondary suite more sensitive to signal plus noise, thus impeding it from transmitting relatively more often. The sensitivity is reduced when the threshold is increased, resulting in increased interference, but also increased secondary utilization. We found 0.8% interference when the threshold is set at 2 dB and 4.7% when set at 8 dB.

VI. CONCLUSION

We have made concurrent spectral measurements at spatially separated sites in the public safety band. These measurements prove useful in examining issues of spectrum sensing for dynamic spectrum access, including primary user signal detection, adjacent channel interference, receiver sensitivity, and policy performance with local and cooperative sensing. It is noted that spectral measurements by low-cost sensors must address some challenges that we have observed in our data, namely differences in sensitivity among sensors, filtering and dynamic range to minimize adjacent channel interference, and methods for correlating measurements between stations. Many issues remain to be investigated, including the spatial correlation among sensing devices at different distances and how to choose an appropriate threshold that distinguishes low-power activities from noise. To choose a threshold, we observe that the consistency of the observed channel occupation over a range of thresholds and a sudden drop as good characteristics to determine an appropriate threshold.

A policy based upon a sensing threshold, a sensing period, and a keep-off time has been evaluated using the measured data to compare secondary utilization with interference providing insight into selection of policy parameters. Selection of a threshold for the detection of primary user signals impacts both secondary utilization and the occurrence of interference to primary users. With the critical importance of the public safety band, it is recommended that a relatively low threshold be chosen so as to minimize interference on the primary users, yet afford a significant secondary throughput in this often sparsely used band. Choice of sensing period is shown to affect secondary utilization in that frequent sensing reduces utilization due to the need to pause transmissions during sensing. Infrequent sensing causes the secondary user to interfere with the primary user for pro-longed periods before detecting the presence of the primary user and abandoning the channel. Peak utilization is found in the middle ground. The occurrence of interference is related to the statistics of the primary user signal. When the primary

signal exhibits highly bursty activity, policy parameters must be modified to track the bursts adequately. These characteristics differ from channel-to-channel and thus policy can be made to be channel dependent.

The value of cooperative sensing has been demonstrated by comparing the interference that would result from local sensing only to that which would result from shared spectral information. In one example, using spectral data measured from a public safety network, we have implemented an assumed policy on that data which would improve interference from 7.2% and 8.7% at two sites resulting from local sensing, to 2.1% overall based upon cooperative sensing.

ACKNOWLEDGMENT

Eric Jung and Xin Liu thank Prof. Zhi Ding for many helpful discussions during the work.

REFERENCES

- [1] Mark McHenry. Spectrum white space measurements. presented to New America Foundation BroadBand Forum, June 2003. http://www.newamerica.net/Download_Docs/pdfs/DocFile_185_1.pdf.
- [2] FCC. FCC spectrum policy task force report, ET Docket No. 02-135.
- [3] FCC. Facilitating opportunities for flexible, efficient, and reliable spectrum use employing cognitive radio technologies, notice of proposed rule making and order, FCC 03-322.
- [4] Sai Shankar N, Carlos Cordeiro and Kiran Challapali. Spectrum Agile Radios: Utilization and Sensing Architectures. *IEEE DYSpan 2005*, Washington DC, USA.
- [5] X. Liu and S. Shankar. Power-sensing based opportunistic access. *ACM Mobile Networks (MONET)*, 2006.
- [6] A. Sahai, N. Hoven, S. M. Mishra, and R. Tandra, "Fundamental tradeoffs in robust spectrum sensing for opportunistic frequency reuse", Tech Report, also submitted to *IEEE JSAC* in 2006.
- [7] Xiangpeng Jing, Siun-Chuon Mau, Dipankar Raychaudhuri and Robert Matyas, "Reactive Cognitive Radio Algorithms for Co-Existence between IEEE 802.11b and 802.16a Networks," *Proceedings of IEEE Globecom*, St. Louis, MO, Nov. 28-Dec. 2, 2005.
- [8] M. Buddhikot, P. Kolodzy, S. Miller, K. Ryan and J. Evans, DIMSUMnet: New Directions in Wireless Networking Using Coordinated Dynamic Spectrum Access, Position Paper in *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (IEEE WoWMoM 2005)*, Taromina/Giardini Naxos, Italy, Jun 2005.
- [9] Nilay Shah, T. Kamakaris, Uf Turelli, and Milind M. Buddhikot, "Wideband Spectrum Sensing Probe for Distributed Measurements in Cellular Band," *Proceedings of First International Workshop on Technology and Spectrum for Accessing Spectrum (TAPAS2006)*, August 5, 2006, Boston, MA.
- [10] Chandrasekharan Raman, Roy Yates and Narayan B. Mandayam, Scheduling Variable Rate Links via a Spectrum Server, *IEEE DySpan 2005*, November 2005.
- [11] T. L. Doumi, Spectrum considerations for public safety in the United States, *Communications Magazine*, IEEE Volume 44, Issue 1, Jan. 2006