

Fundamental Design Tradeoffs in Cognitive Radio Systems

Anant Sahai, Rahul Tandra, Shridhar Mubaraq Mishra, Niels Hoven

School of Electrical Engineering and Computer Science
University of California, Berkeley, California 94704
Email: {sahai, tandra, smm, nhoven}@eecs.berkeley.edu

Abstract—Under the current system of spectrum allocation, rigid partitioning has resulted in vastly underutilized spectrum bands, even in urban locales. Cognitive Radios have been proposed as a way to reuse this underutilized spectrum in an opportunistic manner. To achieve this reuse while guaranteeing non-interference with the primary user, Cognitive Radios must detect very weak primary signals. However, uncertainties in the noise+interference impose a limit on how low of a primary signal can be robustly detected.

In this paper, we show that the presence/absence of possible interference from other opportunistic spectrum users represents a major component of the uncertainty limiting the ability of a Cognitive Radio network to reclaim a band for its use. Coordination among nearby Cognitive Radios is required to control this uncertainty. While this coordination can take a form similar to a traditional MAC protocol for data communication, its role is different in that it aims to *reduce the uncertainty about interference* rather than just reducing the interference itself.

We show how the degree of coordination required can vary based on the coherence times and bandwidths involved, as well as the complexity of the sensors themselves. The simplest sensing strategies end up needing the most coordination, while more complex strategies involving adaptive coherent processing and interference prediction can be individually more robust and thereby reduce the need for coordination across different networks. We also show the existence of a *coordination radius wall* which limits secondary user densities that can be supported irrespective of coordination involved. Furthermore, local cooperation among Cognitive radios for collective decision making can reduce the fading margins we need to budget for. This cooperation benefits from increased secondary user densities and hence induces a minima in the power-coordination tradeoff.

I. INTRODUCTION

Spectrum is considered to be a scarce resource since all frequencies below 3GHz have been completely allocated to specific uses. For example, the Federal Communication Commission's (FCC) frequency allocation chart indicates multiple allocations over all of the frequency bands [1]. However, actual measurements show that most of the allocated spectrum is vastly underutilized at any specific location and time [2]. The FCC has already released a Notice of Proposed Rule Making exploring the operation of unlicensed devices on spatially/temporally "unused" television broadcast bands [3]. These unlicensed devices are deemed to be secondary users of the spectrum since they use the frequency bands only if the official user of the spectrum (primary system) is not using them.

The highest priority constraint for opportunistic secondary devices is maintaining a guarantee of non-interference to the primary system. Given the density of interferers' transmissions, we can determine how large an area we must exclude secondary users from to guarantee an acceptable level of interference to the primary receiver. A secondary user inside this range must be able to robustly detect a primary signal and cease transmission when it is present.

However, if a sense-before-talk network cannot robustly distinguish between the transmissions of the primary system and those of secondary systems, it will fail to capitalize on situations where the spectrum is available and already being used by another secondary system. One can imagine a "finders keepers" scenario in which the first secondary system to discover a band keeps it to itself because the other systems hear the secondary transmissions and assume the primary signal is present. This is unfair; we would like spectrum usage to be shared among secondary transmitters.

Furthermore, when spectrum-agile nodes mistakenly detect a primary in a given frequency band, they switch to another. This effectively forces the secondary systems to orthogonalize. Because this orthogonalization can lead to a significant loss of spectrum utilization, it should be driven by the requirements of the data MAC [4], not the sensing protocol. The data MAC has a much higher tolerance for interference than does the sensing protocol, which has the additional constraint of non-interference to the primary system.

We show that the presence/absence of possible interference from other opportunistic spectrum users represents a major component of the uncertainties limiting the ability of a sensor network to reclaim a band for its use. We begin by describing and then quantifying the fundamental issues governing opportunistic system design (Section II). We follow this by reviewing the performance of radiometric and coherent detectors in see Section III and quantify the performance impact of uncertainty and multipath on these detectors. We also explore various sources of uncertainty and the limits of interference estimation in reducing this uncertainty.

In Section IV we describe the power-coordination tradeoffs necessary to reduce this uncertainty emphasizing the conditions under which the system can encounter a coordination radius wall.

Finally, in Section V we explore the impact of cooperating

users on fading margins and quantify the power-coordination tradeoffs with cooperating users.

II. COGNITIVE RADIO SYSTEMS: THE DESIGN PROBLEM EXPLAINED

Much of the following material is addressed in greater detail in our previous papers. We review it here partially for completeness and partially to emphasize issues whose fundamental importance has only become clear since our last publication.

A. Fundamental tradeoffs

A number of concerns must be alleviated before regulators will allow Cognitive Radios to transmit opportunistically in already allocated bands. The most important constraint is that of non-interference to the primary receivers. While a traditional link budget analysis reflects the impact of these factors on a system's minimum transmit power, under the opportunistic regime the constraint is placed upon a system's maximum transmit power. Alternatively, if a system's transmit power is fixed, it dictates how far away the system must be before it is allowed to transmit.

Figure 1 illustrates the fundamental tradeoffs between a primary receiver's margin of protection and a single opportunistic user's transmission power [5]. The large antenna at the center represents a primary system's transmitter. These primary users may be providing socially important services, or they might simply be legacy systems that are unable to change. The outer circle (1a) represents the boundary of decodability for a single-user system. That is, in the absence of all interference, a user within the outer circle would be able to decode a signal from the transmitter, while a user outside the circle would not.




Fig. 1. Weaker secondary users can transmit closer to the primary receiver (r_p -protected radius, r_{dec} -decodability radius, r_n - no talk radius i.e. radius around the Primary transmitter where that no secondary transmitter can transmit)

The inner shaded circle represents the “protected radius” where we guarantee decodability to primary receivers. The more we shrink the bound of the protected region inside the decodability region, the smaller the necessary no-talk zones become. Similarly, the power of the secondary user's transmissions is important. If they are “mice,” (Figure 1a) who squeak softly, then the no-talk radius can be much smaller. If they are “lions,” (Figure 1b) roaring with high power transmissions, the radius of the no-talk zones will become much larger.

To ensure that cognitive radios are low cost devices, we assume they do not contain any GPS systems, or other hardware to determine location. Nor is a system designer likely to want to program location information into each node individually before installing them carefully in a predetermined deployment pattern. Finally, the service areas of the primary systems might be unknown or time-varying. Hence it is expected that secondary radios, after detecting the primary signal, will estimate the signal's strength and use it as a proxy for distance from the primary transmitter.

For a system with a Δ dB drop between the primary transmitter and its decodability radius, a μ dB margin of protection¹ for primary receivers, and attenuation $g_{21}(r)$ ($r^{-\alpha_{21}}$, for example) between a primary transmitter and secondary receiver, we can express how much more sensitive (ψ dB) secondary detectors must be than the primary receivers [6], [7]. The formula illustrates the idea that we must think in terms of distances, but calculate at the level of signal strengths.

$$\begin{aligned} \psi + \Delta &\geq -10 \log_{10} [g_{11}(r_2)] \\ &= -10 \log_{10} [g_{11}(r_p + (r_n - r_p))] \\ &= -10 \log_{10} \left[g_{11}^{-1} \left(10^{\frac{\mu - \Delta}{10}} \right) + g_{21}^{-1} \left(\frac{\sigma^2 \cdot (10^{\frac{\mu}{10}} - 1)}{P_2} \right) \right] \end{aligned}$$

In the above equation g_{11} is the attenuation between the primary transmitter and the primary receiver while g_{21} is the attenuation between the secondary transmitter and the primary receiver.

B. Multiple secondary users

Further complications arise in the case of multiple secondary users. A primary receiver now experiences aggregate interference from all transmitting secondary devices. The aggregate interference from multiple secondary users transmitting simultaneously at the same power (assuming signal attenuation as $d^{-\alpha_{21}}$) behaves like a single transmitter with a slower rate of attenuation ($d^{-\alpha_{21}+2}$) [6], [7].

$$I_{aggregate}(r) = K(\alpha_{21})D(r_n - r)^{-\alpha_{21}+2}$$

In the above equation, $I_{aggregate}(r)$ is the aggregate interference from all secondary transmitters at a distance of r from the primary transmitter (our interest is the interference at a distance r_p), $K(\alpha_{21})$ is a integration constant, D is the power density (in units of dBW/m^2) and $-\alpha_{21}$ is the attenuation coefficient (assuming a polynomial decay model) between secondary transmitter and the primary receiver.

An alternative view is that the primary receivers have a certain margin of interference that they can tolerate, and this margin must be divided among the secondary users. Nodes further from the primary system can transmit at greater powers, but only if nodes closer in transmit more quietly. Therefore,

¹The margin of protection is how much of an increase in the interference+noise floor that the primary system can tolerate without losing service to the protected receivers. Not all potential receivers will be deemed protected receivers otherwise no such margin μ exists.

heterogeneous transmit powers can be accommodated, but policy decisions are necessary to fix the tradeoffs.

C. Shadowing/Fading

Even if there are only a few primary receivers, secondary users must stay out of the space that is the union of all possible no-talk zones (Figure 2a) (secondary transmitters can exist in the external grey zone). We note that in the hypothetical example, the prohibited region for secondary users has already extended beyond the decodability region.




Fig. 2. Uncertainty in the locations of primary receivers results in a larger no-talk zone. Uncertainty from shadowing increase the zone even further.

If we take fading into account (Figure 2b), the prohibited region continues to grow. Shadowing and multipath fading can cause the secondary user in the notalk zone to see a very low power primary transmission. As a result, a secondary has no way to tell if it is safely outside the protected region (ST2) or in the global quiet zone but behind a building (ST1). To avoid secondary users in a local shadow interfering with unshadowed primary receiver, the no-talk zone must be pushed out even further. The primary signal power here is low enough (details in Section V) for users to assume that it is unlikely that they are just in a local shadow. The possibility of β dB of shadowing translates to a pure shift of β in the required sensitivity of the opportunistic network.

$$\psi + \Delta \geq -10 \log_{10} \left[g_{11} \left(g_{11}^{-1} \left(10^{\frac{\mu - \Delta}{10}} \right) + g_{21}^{-1} \left(\frac{\sigma^2 \cdot (10^{\frac{\mu}{10}} - 1)}{P_2} \right) \right) \right] + \beta$$

If deeply faded signals must be detected, secondary devices must be extremely sensitive to coexist with the primary system.

D. Limits on sensing

As explained in the Section III a single radio cannot robustly detect signals below a certain SNR (called the SNR threshold). This problem is particularly severe for radiometric detectors but also limits the performance of coherent detectors due to coherence time constraints.

This fundamental limitation on sensing arises because at low SNR it is difficult to be robust to even small system uncertainties. In particular, we consider uncertainties in the noise plus interference and try to understand its effect on detection. For example, any detector suffers from device level uncertainties due to non-linearity of various components,

non-uniform, time-varying thermal noise etc. We can model this effect by assuming that the device level noise power is uncertain and can take any value in $[\frac{1}{\rho}\sigma^2, \rho\sigma^2]$, where ρ parameterizes the amount of uncertainty in the system and σ^2 is a nominal value of thermal noise power. This uncertainty has been experimentally observed in [?].

However, the dominant source of uncertainty for a Cognitive Radio network is *potential interference* from transmissions of other opportunistic devices communicating within the same band. The uncertainty in this case arises because the sensing Cognitive Radio is unable to ascertain how many of the surrounding Cognitive radios are transmitting. In general the noise uncertainty set looks like $[\sigma_{low}^2, \sigma_{high}^2]$ where σ_{low}^2 corresponds to the case where none of the surrounding radios are transmitting while σ_{high}^2 corresponds to the case where all the radios are transmitting.

While device noise uncertainty cannot be reduced, uncertainty from other secondaries can be reduced by implementing a coordination radius around the sensing radio. All radios within this radius need to cease transmission in order to reduce the uncertainty to acceptable levels where acceptable depends on the minimum primary signal level that we intend to detect. Hence the procedure for designing the system works as follows: We start with a protection margin (μ dB) and select a secondary power density that we intend to support. This in turn yields a no-talk radius. The no talk radius and the fading margin dictates the minimum power level at which the primary has to be detected. This level limits the uncertainty that we can deal with. The uncertainty requirements mandate the coordination radius. These power-coordination tradeoffs are explored in Section IV but first we examine the impact of uncertainty on a single radio.

III. SINGLE RADIO DETECTOR: LIMITS ON ROBUST SENSING

To meet the constraint of non-interference to the primary receivers, the sensing cognitive radio must first be able to detect the presence or absence of the primary signal. For this the cognitive radio must be able to sense the primary signal in the ‘band of interest’, as shown in figure 3. There are several possible decision strategies that can be used for sensing – energy detection, coherent detection, and many others somewhere in the middle. In this section we consider both the energy detector and the coherent detector and review the fundamental limits on robust detection due to uncertainties.

Spectrum picture




Fig. 3. Big question: sensing the primary

Sampling the ‘band of interest’ at Nyquist, the detection problem can be formulated as a binary hypothesis testing problem, where the aim is to distinguish the following hypotheses:

$$\begin{aligned} \mathcal{H}_0 : Y[n] &= W[n] & n = 1, \dots, N \\ \mathcal{H}_1 : Y[n] &= X[n] + W[n] & n = 1, \dots, N \end{aligned} \quad (1)$$

Here $X[n]$ are i.i.d signal samples and $W[n]$ i.i.d noise samples. We are interested in decision strategies that achieve a given target probability of false alarm, P_{FA} , and probability of missed detection, P_{MD} . Complete analysis of this detection problem under various performance measures is done in [8], [9]. In particular they give fundamental limitations on detection due to uncertainties. We summarize the main results of [8], [9] here to make this paper self-contained.

A. Radiometric detection

The radiometer is a simple detector which senses for the presence/absence of the primary signal based on the energy of the received signal samples. The test statistic for this detector is $\frac{1}{N} \sum_{n=1}^N Y[n]^2$ which is compared to a threshold to distinguish between the two binary hypotheses in (1). We now see the effect of noise uncertainty on energy detection. To start, we consider only device level noise uncertainty. This can be modeled by assuming that the noise variance can take any value in $[\frac{1}{\rho}\sigma^2, \rho\sigma^2]$. In this case, it is clear that the radiometer will not be able to detect the signal if the signal strength is less than $(\frac{\rho^2-1}{\rho})\sigma^2$, i.e., the radiometer fails to detect below an SNR threshold. This effect can be understood from figure 4. The shaded area in the Figure represents the uncertainty in the noise power. It is clear from the Figure that if the test statistic falls within the shaded region, there is no way to distinguish between the two hypotheses. Hence the radiometer hits an SNR wall, and is non-robust to simple device level uncertainties.




Fig. 4. Understanding device level noise uncertainty in the case of a radiometer

However, as discussed in Section II-D the dominant source of uncertainty for a cognitive radio network is potential transmissions of other opportunistic devices communicating within the same band. In general the noise uncertainty set is of the form $[\sigma_{low}^2, \sigma_{high}^2]$ and the radiometer fails if the primary signal strength, $P \leq (\sigma_{high}^2 - \sigma_{low}^2)$. This clearly

shows that interference from other opportunistic devices makes the radiometer highly non-robust.

Furthermore, it has been shown in [?], [?] that similar limits on robust detection hold as long as the primary signal has no deterministic mean and is sampled at Nyquist. In this case for every detector, there always exists an SNR threshold below which detection is impossible.

B. Coherent detection

Since most primary signals contain deterministic pilot tones we can try to coherently detect the primary’s pilot tone. We now analyze the coherent detector in detail and evaluate the possible coherent processing gains.

1) *Hypothesis testing model: Coherent detection:* We can again model the problem of detecting the primary’s pilot tone as a binary hypothesis testing problem:

$$\begin{aligned} \mathcal{H}_0 : Y[n] &= W[n] \\ \mathcal{H}_1 : Y[n] &= \sqrt{\theta}X_p[n] + \sqrt{(1-\theta)}X[n] + W[n] \end{aligned} \quad (2)$$

where the signal samples $X[n]$ are assumed to be white, or orthogonal to the pilot, and noise samples $W[n]$ are assumed to be white. $X_p[n]$ is a known pilot tone with θ being the total fraction of energy allocated to the pilot tone. The test statistic for the coherent detector is

$$T(\mathbf{y}) = \frac{1}{N} \sum_{n=1}^N Y[n] \hat{X}_p[n] \quad (3)$$

where \hat{x}_p is a unit vector in the direction of the pilot tone. It is known that in the simple case of completely known noise statistics and no uncertainties in the system the optimal detector for (2) is a matched filter. It achieves a dwell time $N \approx [\mathcal{Q}^{-1}(P_D) - \mathcal{Q}^{-1}(P_{FA})]^2 \theta^{-1} SNR^{-1}$, [9].

However, we are interested in the case when there is uncertainty in the interference. We model noise uncertainty by assuming that the noise variance lies in $[\sigma_{low}^2, \sigma_{high}^2]$. From (3) one can view the test statistic as obtained by projecting the received signal onto the pilot direction. This averages out the noise in other orthogonal directions and the new noise uncertainty set is $[\frac{\sigma_{low}^2}{N}, \frac{\sigma_{high}^2}{N}]$, which shrinks to the null set as $N \rightarrow \infty$. So, by dwelling long enough, the primary signal strength becomes greater than the uncertainty in the noise and hence there is no SNR wall. See [?] for a more detailed analysis of the coherent detector. This clearly shows that the coherent detector is robust to noise level uncertainties. The only penalty due to noise uncertainty is increase in the dwell time, $N \approx [\mathcal{Q}^{-1}(P_D) - \mathcal{Q}^{-1}(P_{FA})]^2 \theta^{-1} \left(\frac{P}{\sigma_{high}^2}\right)^{-1}$.

The above result implies that the gains from coherent processing can overcome the uncertainty from interference and hence local cooperation between different systems is unnecessary. But, the luxury of non-cooperation exists only if we can get infinitely large gains from coherent processing. Unfortunately, in reality the gains from coherent processing are limited!

C. Limited coherent processing: fundamentally challenging uncertainty

Incorporating fading into our model, the detection problem is to distinguish between the following hypotheses:

$$\begin{aligned}\mathcal{H}_0 : Y[n] &= W[n] \\ \mathcal{H}_1 : Y[n] &= \sum_{l=0}^{L-1} h_l[n] \tilde{X}[n-l] + W[n]\end{aligned}$$

where $\tilde{X}[n] = \sqrt{\theta}X_p[n] + \sqrt{(1-\theta)}X[n]$, and $h_l[n]$ are the multipath fading coefficients. As before, we also assume that the noise level is uncertain, i.e., the noise variance lies in $[\sigma_{low}^2, \sigma_{high}^2]$.

In this case it is clear that we cannot reap the gains of coherent signal processing forever. As soon as the channel taps assume independent realizations, we can no longer gain from coherent signal processing². So, the detector performance depends on the channel's potentially shortest³ coherence time T_c . Here channel coherence time is defined as the time for which the channel taps remain sufficiently constant to allow matched filtering (see [11]). For simplicity let us assume that the channel coherence time is known at the sensor.

It can be shown that the optimal thing to do in this case is to combine the signal coherently in each coherence block and then use the radiometer over multiple coherent blocks (see [?] for details). Specifically, the test statistic is given by

$$T(\mathbf{y}) = \frac{1}{N} \sum_{n=0}^{N-1} \left[\frac{1}{\sqrt{N_c}} \sum_{k=1}^{N_c} Y[n] \hat{X}_p[nN_c + k] \right]^2$$

where N_c is the number of samples in each coherent block. This detector can be visualized as a combination of two detectors. First, the signal is coherently combined within each coherence time T_c . Due to this the signal strength gets boosted by a factor N_c while the noise power is unchanged. Secondly, the new boosted signal is detected at the receiver by passing it through a radiometer.

Since we are reduced to an energy detector, our final detector is still non-robust to noise level uncertainties. However, the robustness is improved and the SNR wall is lowered by a factor of $10 \log_{10} N_c$ dB. As far as the dwell times are considered, the performance of this detector lies strictly in between the utopian matched filter and the radiometer.

Coherent processing gains could be limited because of complexity reasons too. The clock-instability of the cognitive nodes also imposes a limit on the coherent processing time. For instance, consider the following situation: Assume that there is 1000 Hz of frequency uncertainty and we are doing coherent processing over 1 ms. Since the pilot frequency is uncertain, we need to search over different frequency bins to

locate the pilot tone. Let's assume that we need to search over 4 bins to achieve a target probability of false alarm. However, if we decide to do coherent processing over 10 ms, we now need to search over 40 frequency bins. Since searching over frequency bins involves computing an FFT, it is clear that the coherent processing time is limited by the complexity of the node.

These results clearly show that both the radiometer and the coherent detector are non-robust to uncertainty in the noise plus interference level. It is also clear that the dominating effect is the uncertainty in the interference level and the SNR wall depends on the net uncertainty in noise plus interference, $(\sigma_{high}^2 - \sigma_{low}^2)$. This in turn directly affects the mandatory coordination radius required for sensing: larger noise uncertainty imposes a larger coordination requirement while sensing. Therefore, any strategy to reduce the uncertainty in the noise plus interference will improve the robustness of sensing and also the coordination radius. Hence, we first try to robustly estimate the interference level.

D. Interference estimation

In order to reduce the mandatory coordination radius, it is clear that we need to estimate the interference from the secondary transmissions.




Fig. 5. Guard bands measurements can help estimate interference

The frequency spectrum is broken up into different chunks of bandwidth and is allocated to primary users by the FCC. However, for practical reasons small guard bands (see fig. 5) between these spectrum chunks are left unallocated. We can use these guard bands to estimate the interference. For instance consider the scenario in figure 5. By measuring the energy in the guard bands adjacent to the 'band of interest' we can get an estimate of the noise plus interference (see figure 6). This can possibly reduce uncertainty in the level of interference, which in turn lowers the coordination radius and also improves robustness. However, the actual gain from interference estimation depends on the quality of our estimate.

The accuracy of our estimate depends on several factors. The guard bands have leakages from adjacent bands which affect our interference estimate. Also, the estimate depends on the frequency selectivity of the secondary transmissions. For example, some secondary systems might be transmitting OFDM packets, in which case the interference varies with the frequency of the subcarriers actually used. So, the measurement in the guard band may or may not be a good estimate of the interference in the center of the primary band.

²Here we are making the assumption that the channel taps cannot be tracked at the receiver. This assumption is valid since we are operating at low SNR and we are not even sure if the primary signal is present or absent.

³We can not count on the typical coherence time since that would mean that many nodes are not actually that sensitive and their decisions can not be trusted. This issue of non-trust is addressed in [10].




Fig. 6. Interference estimation via guard band measurements

However, if the cognitive radio is trying to detect the primary's pilot tone, then the estimation error can be significantly reduced. This is because we can estimate the interference level by measuring the energy at frequencies near the pilot frequency rather than in far away guard bands. This way we can better estimate the noise plus interference level for setting the coherent detection thresholds.




Fig. 7. In band interference estimation

It is clear that the quality of the estimate depends on the correlation of the interference across frequency. The reference location must be separated from the pilot frequency by at least the Doppler spread, D_s , of the primary-secondary channel. This is to ensure that the interference estimate has very little contribution from the energy of the pilot tone. Furthermore, the reference location must also be farther than the uncertainty in the clock frequency. In most cases the frequency uncertainty is larger than the Doppler spread. On the other hand, the reference location must be within the same coherence bandwidth, W_c , (of the wireless channel from secondary-secondary, i.e., the channel between the cognitive radios) as the pilot frequency to ensure that the interference does not change too much (see fig. 7). So there is an apparent tension between these requirements. However, since the cognitive radios are not spaced too far, the delay spread, T_d , of the channel is probably very low, and hence $W_c \approx \frac{1}{T_d} \gg \max(D_s, \text{frequency uncertainty})$. Therefore, the interference estimate is potentially very accurate. When we combine this gain with coherent processing gains, the mandatory coordination radius is reduced significantly. See Section IV for numerical results.

IV. CONCEPTUALIZING SECONDARY-TO-SECONDARY INTERFERENCE

As seen in the previous section, uncertainty limits the performance of radiometer as well as coherent detectors. To robustly detect the primary signal, this uncertainty must be reduced. Either all the secondary systems must be constrained to operate at very low power densities or there must be coordination among different systems. We now propose the following form of local coordination: systems within a certain radius of the sensing node must remain quiet for the duration of sensing. The tradeoff between the sensing radius and the secondary power density is developed in the next section.

A. Power-Coordination tradeoffs

We propose the existence of a mandated “sensing MAC” among systems which ensures that whenever a particular cognitive node is sensing for the primary signal, all nodes within a certain coordination radius, r_s , refrain from transmitting. This sensing MAC must ensure that the uncertainty due to the interference from secondary transmissions is tolerable.

To arrive at the required power-cooperation tradeoff, we first start with a given secondary power density D . This power density contributes to some aggregate interference, $I_{aggregate}(r)$ at the primary receivers located at a distance of r from the primary transmitter. We must ensure that all primary receivers within the protected radius are still able to decode their primary transmissions. Under this constraint it is possible to calculate a lower bound on the no-talk radius:

$$r_n \geq r_p + \left[\frac{K(\alpha_{21})D}{(10^{\frac{\beta}{10}} - 1)\sigma^2} \right]^{\frac{1}{\alpha_{21}-2}} \quad (4)$$

The sensors at the edge of the no-talk radius r_n receive a primary signal of power $P_n = P_{primary}r_n^{-\alpha_{12}}10^{\frac{\beta}{10}}$. Here we account for β dB of possible fading and shadowing. For the radio to be able to detect the primary signal, P_n must be greater than the uncertainty in the noise plus interference. The uncertainty at the sensor can be divided into two categories: device level uncertainty and uncertainty due to other secondary transmissions. The device level uncertainty can be accounted for by assuming that the noise variance lies in $[\frac{1}{\rho}\sigma^2, \rho\sigma^2]$ for some suitable $\rho \geq 1$. Since we do not know the number of actual interferers, the total interference can lie anywhere in $[0, I_{max}]$, where I_{max} is the maximum possible interference from secondaries outside the shut-up radius r_s and is given by

$$I_{max} = D \frac{2\pi}{\alpha_{22} - 2} r_s^{-\alpha_{22} + 2}, \quad (5)$$

Therefore, the actual noise plus interference can lie in $[\frac{1}{\rho}\sigma^2, \rho\sigma^2 + I_{max}]$. So in order to overcome the SNR wall we must have

$$P_n \geq I_{max} + \frac{\rho^2 - 1}{\rho} \quad (6)$$

From (4), (6) and (5) we get the required tradeoff between D and r_s :

$$r_s \geq \left[\frac{D \left(\frac{2\pi}{\alpha_{22}-2} \right)}{P_n - \frac{\alpha^2-1}{\alpha}} \right]^{\frac{1}{\alpha_{22}-2}} \quad (7)$$

B. Power-Coordination tradeoffs: Coherent case

Now, we re-derive the power-cooperation tradeoff given in (7) for the coherent detector also taking into account in-band interference estimation. Let λ denote the percent interference estimation error (see fig. 6). The new tradeoff is given by

$$r_s \geq \left[\frac{D \left(\frac{2\pi}{\alpha_{22}-2} \right) \lambda}{z\theta P_n - \frac{\alpha^2-1}{\alpha}} \right]^{\frac{1}{\alpha_{22}-2}} \quad (8)$$

where z is the power gain obtained by coherent signal processing with a single coherent block. The above tradeoff is very similar to the one in (7). The λ in the numerator is because the uncertainty in the interference is reduced from I_{max} to λI_{max} . Also, the effective signal power at the no-talk region is $z\theta P_n$ because the sensors try to detect the primary pilot tone and the pilot strength is boosted due to coherent signal processing within each channel coherence block.

C. Power-Coordination tradeoffs: Results

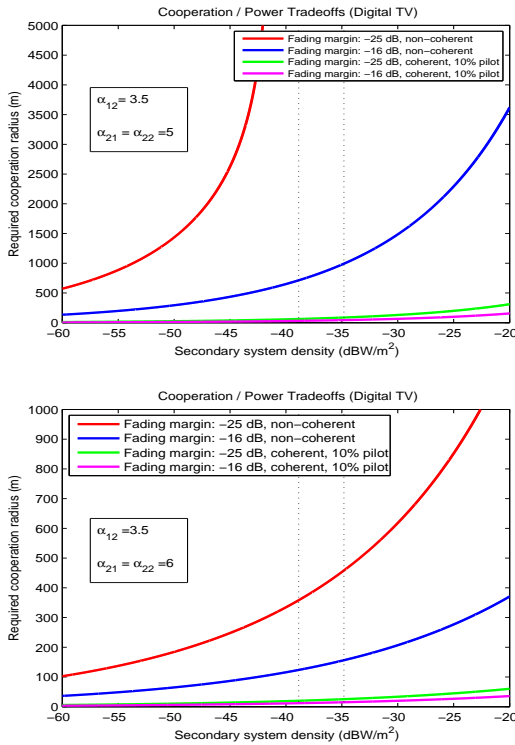


Fig. 8. The coordination radius r_s is plotted as a function of the secondary power density. The top plot corresponds to a secondary power decay of r^{-5} and the bottom one corresponds to a secondary power decay of r^{-6} .

Figure 8 plots the power-cooperation tradeoff for different values of fading margins, power decay models and detectors. The two fading margins (25dB and 16dB) used in the

graph correspond to 90% and 50% probability of detection respectively. Results from the top Figure have been plotted using an r^{-5} power decay rule for secondary transmissions while the lower figure is generated using an r^{-6} decay model for secondary transmissions. This figure clearly illustrates that coordination is essential to mitigate the effect of fading in an energy detector. The coordination radius for coherent detectors is much smaller due to the SNR gains achieved from coherent processing during the coherence time. Even for reasonable power density levels, the radiometer coordination radius r_s is of the order of hundreds of meters (or a few kilometers) for the r^{-5} decay rule, which is ridiculously high. In contrast to this, the coordination radius is much smaller for the r^{-6} decay rule. This shows that the power-coordination tradeoff obtained is very sensitive to assumptions on the power decay models. However, in both models the required coordination radius is unreasonably high.

Another interesting trend is the presence of a coordination radius wall which is visible for the red curve in the upper graph. The coordination wall exists when power densities beyond a certain limit cannot be robustly supported. At these densities, the no-talk radius is too large which causes the minimum detectable signal to be small. As seen in Section III-A, to detect such a small signal, the uncertainty has to be very low. The irreducible device level uncertainty now limits the detection and no coordination can help in this case. The coordination wall exists for all detectors and fading margins though it is not visible for the power densities depicted in the graphs.

V. THE ROLE OF INTRA-SYSTEM COOPERATION

In Section II-C we have shown that the fading margin β can adversely effect the power-coordination tradeoff. Fortunately, the fading margins can be reduced by using multiple radios as shown in [10].

Without cooperation, an opportunistic system must increase the sensitivity of its detection algorithm by an amount equal to that of the deepest shadow/fade under which it must detect. That is, while an unshadowed system might only need to detect a 0dB signal to avoid interfering, it would need to detect a -56dB signal to account for the possibility of a 56 dB shadow/fade.

Device sensitivity becomes a serious issue if there are K secondary systems. If the primary system is willing to tolerate a probability of harmful interference P_{HI} , the probability of missed detection for one specific secondary system must now be less than P_{HI}/K [10].

Even with just a few co-located Cognitive networks, requiring a probability of missed detection on the order of 0.1% or less is unrealistic. First, to achieve this P_{md} , a node must be able to detect signals attenuated by the 0.1% deepest fades, which could be tens of dB under some models. Second, no model accurately describes such low probability fading behavior.

Fortunately, independent measurements of the faded signal drastically improve the individual detection requirements [10].

With M independent measurements, the required probability of missed detection increases to $\sqrt[M]{P_{HI}/K}$. Cooperation among nodes can therefore raise the detection threshold to reasonable levels for a practical system.

Large cooperative gain (which in turn implies small fading margins) requires many users. Increasing users in a given area to achieve these low fading margins also increases the secondary power density. In this case we are assuming that each radio has a fixed power and hence increasing user densities increases the secondary power density. Figure 9 shows the fading margins obtained from increased user densities.




Fig. 9. Fading margin is plotted with increased power densities. The blue curve corresponds to the case where all users encounter independent multipath and shadowing. The red curve corresponds to the case where all radios see independent multipath but the shadowing gains are only from user which are not behind the same obstacle. On an average, with N users in a given area, only \sqrt{N} will contribute to shadowing gains. Furthermore the multipath gains in the red curve have been capped at 0

Factoring density dependent fading margins obtained from Figure 9 into the calculation of the coordination radius of Figure 8, we see an interesting trend where the coordination radius exhibits a minima at certain secondary power densities (See Figure 10). At lower densities, the no-talk radius is small but the fading margin is large. As the user density is increased, the fading margin decreases due to cooperative gains though the no-talk radius increases and hence the power loss due to the decay model increases. These opposing effects induce the minima.

The straight line at 312m corresponds to the cooperation footprint. Radios within this radius of a sensing radio sense together and collaborate to make a collective decision. We would prefer to keep the coordination radius approximately the same as the cooperation footprint so as to make sure that all radios within the system are sensing at the same time and collectively making decisions. If secondary systems are not geographically interpenetrating, this eliminates the need for coordinating across systems. As seen in the top figure in figure 10, achieving this is almost impossible with an energy detector with a secondary power decay of r^{-5} .

Again we observe an interesting phenomenon — a wall on the left i.e. when we include the effects of cooperation, certain *low densities* of users also cannot be supported. This

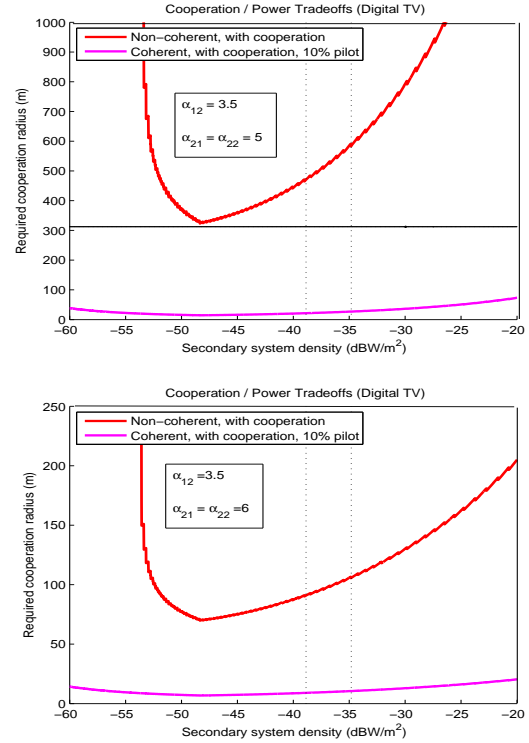


Fig. 10. The coordination radius r_s is plotted as a function of the secondary power density with cooperation among users. The top plot corresponds to a secondary power decay of r^{-5} and the bottom one corresponds to a secondary power decay of r^{-6} .

is because at low densities, the cooperation gain is small (and so a correspondingly large fading margin has to be set) and uncertainties are dominated by radio level uncertainty. In general, both lower and upper coordination walls exist, i.e. too small and too large power densities cannot be supported.

VI. CONCLUSION

The presence/absence of possible interference from other opportunistic spectrum users represents a major component of the uncertainty limiting the ability of a sensor network to reclaim a band for its use. This uncertainty must therefore be reduced through some sort of inter-system coordination. This cooperation has the flavor of a traditional MAC protocol, but a data MAC has a much higher tolerance for interference than does the sensing protocol, which has the additional constraint of non-interference to the primary system. We have shown that the amount of inter-system cooperation required can be quite high and hence have considered the gains possible with more complex signal detectors.

For opportunistic spectrum use, there are three kinds of complexity: complexity of the detector, complexity of intra-system cooperation, and complexity of inter-system coordination. There is a basic conservation of complexity involved in the design of opportunistic systems - one cannot decrease the complexity of one part of a network without compensating for it elsewhere. While standardization of a protocol may eventually make intra-system coordination automatic, the designer

of any opportunistic sensor network must take these tradeoffs into account.

REFERENCES

- [1] NTIA, "U.S. frequency allocations." [Online]. Available: <http://www.ntia.doc.gov/osmhome/allochrt.pdf>
- [2] R. W. Broderson, A. Wolisz, D. Cabric, S. M. Mishra, and D. Willkomm, "White paper: CORVUS: A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum," Tech. Rep., 2004. [Online]. Available: http://bwrc.eecs.berkeley.edu/Research/MCMA/CR-White_paper_final1.pdf
- [3] FCC, "FCC 04-113," May 2004. [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-04-113A1.pdf
- [4] R. Etkin, A. Parekh, and D. Tse, "Spectrum sharing for unlicensed bands," in *IEEE DySPAN 2005*, Baltimore, MD, Nov.8–11 2005.
- [5] A. Sahai, N. Hoven, and R. Tandra, "Some fundamental limits on cognitive radio," in *Forty-second Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2004.
- [6] N. Hoven and A. Sahai, "Power scaling for cognitive radio," in *Proc. of the WirelessCom 05 Symposium on Signal Processing*, Maui, HI, June 13–16 2005.
- [7] N. Hoven, "On the feasibility of cognitive radio," Master's thesis, University of California, Berkeley, 2005.
- [8] R. Tandra and A. Sahai, "Fundamental limits on detection in low SNR under noise uncertainty," in *Proc. of the WirelessCom 05 Symposium on Signal Processing*, Maui, HI, June 13–16 2005.
- [9] R. Tandra, "Fundamental limits on detection in low SNR," Master's thesis, University of California, Berkeley, 2005.
- [10] S. M. Mishra, A. Sahai, and R. W. Broderson, "Cooperative sensing among cognitive radios," in *ICC 2005*, Istanbul, Turkey, June 11–15, 2006.
- [11] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. New York: Cambridge University Press, 2005.