

What is a spectrum hole and what does it take to recognize one?

Rahul Tandra Shridhar Mubaraq Mishra Anant Sahai
 tandra@eecs.berkeley.edu smm@eecs.berkeley.edu sahai@eecs.berkeley.edu
 Dept. of Electrical Engineering and Computer Sciences, U C Berkeley

Abstract—“Spectrum holes” represent the potential opportunities for non-interfering (safe) use of spectrum and can be considered as multidimensional regions within frequency, time, and space. The main challenge for secondary radio systems is to be able to robustly sense when they are within such a spectrum hole. To allow a unified discussion of the core issues in spectrum sensing, the “Weighted Probability of Area Recovered (WPAR)” metric is introduced to measure the performance of a sensing strategy and the “Fear of Harmful Interference” F_{HI} metric is introduced to measure its safety. These metrics explicitly consider the impact of asymmetric uncertainties (and misaligned incentives) in the system model. Furthermore, they allow a meaningful comparison of diverse approaches to spectrum sensing unlike the traditional triad of sensitivity, probability of false-alarm P_{FA} , and probability of missed detection P_{MD} . These new metrics are used to show that fading uncertainty forces the WPAR performance of single-radio sensing algorithms to be very low for small values of F_{HI} , even for ideal detectors. Cooperative sensing algorithms enable a much higher WPAR, but only if users are guaranteed to experience independent fading. Finally, in-the-field calibration for wideband (but uncertain) environment variables (e.g. interference and shadowing) can robustly guarantee safety (low F_{HI}) even in the face of potentially correlated users without sacrificing WPAR.

I. INTRODUCTION

Wireless systems deliver real value to their users, but require radio spectrum to operate. The use of a band of spectrum by one system in the vicinity of a second system’s receiver (tuned to the same band) will generally degrade the performance of that second system if the total interference exceeds a critical value.¹ Therefore, spectrum is in principle a potentially scarce resource. Indeed, across the planet, spectrum is regulated so that most bands are allocated exclusively to a particular service, often with only a single system licensed to use that band in any given location. It is generally illegal to transmit without an explicit license. It is the *fear* of harmful interference that drives this policy of prior restraint.

This approach has been largely successful in avoiding interference, but in practice it does so at the expense of overall utilization. *Most bands in most places are underused most of*

the time [3]–[5]. A band of spectrum can be considered under-used if it can accommodate secondary transmissions without harming the operation of the primary user of the band.² The region of space-time-frequency in which a particular secondary use is possible is called a ‘spectrum hole.’ Spectrum holes are defined and discussed further in Section II.

Upon reflection, spectrum holes are a natural consequence of the gap between the distinct scales at which regulation and use occur — *just as a vase can be filled with rocks and still have plenty of room for sand*. Spectrum regulatory agencies perform allocations that are valid for multiple years/decades and over spatial extents that are hundreds of miles across. This is despite the fact that useful spectrum use could occur even over a few milliseconds and in a manner that is localized around transmitter-receiver pairs only tens of meters apart.

Why then do not regulatory agencies simply adjust their regulatory granularity to deal with scales closer to those of actual use? If a *static* approach to spectrum access is assumed wherein devices and wireless systems are inherently tied to particular bands and the regulator acts by certifying devices and systems before they are put into service, then the regulatory granularity is lower-bounded by the natural lifespans of wireless systems and the mobility of the devices. The lifespan of a wireless system is governed by the business models for the service — the system has to operate for long enough to result in a positive return on the infrastructure investments. The lifetime might differ wildly from one application to another³ — and thus by Moore’s law, the technical sophistication of wireless systems can and will differ greatly from each other. The freedom of innovation and movement for the users of one system translates into uncertainty for the operators of another. *The unknown is feared if it can affect you*. To reduce this fear of harmful interference, the interaction must be precluded by ensuring that different users are in different bands even after they have physically moved.

Yet the overall demand mix for different applications/services is almost certain to be different from one location to another, and so in a world of heterogeneous wireless

¹The performance degradation with increased interference can be gradual in the case of analog systems or catastrophic in the case of digital systems. While the critical value of total interference is therefore relatively unambiguous for digital receivers, a subjective judgment of “minimally acceptable quality” is required for analog systems. In the literature, the critical value of total interference is called the “interference temperature limit” [1], [2]. The terminology itself is meant to suggest that interference can be considered to be like additional thermal noise.

²Using the language of interference temperature, underutilization is said to exist whenever the actual interference temperature at a location has not yet reached the specified interference temperature limit [1], [2]. However, it turns out that interference temperature alone is not enough to understand the concept of a spectrum hole [6]–[8].

³Compare the longevity of analog television to the different cellular or Wireless Local Area Network (WLAN) standards that have come and gone within the same time period.

services and static allocations, waste is seemingly unavoidable. This also precludes otherwise brilliant approaches (see e.g. [9], [10]) that design transmissions so that the interference at receivers is aligned roughly orthogonal to their desired signals. Such an approach is not practical for heterogeneous services because it requires the potentially interacting systems to jointly coordinate their transmissions.

Bridging this gap and filling in spectrum holes requires a *dynamic* approach to spectrum access. Wireless systems must determine where the holes exist and reconfigure to take advantage of these opportunities. Regulation shifts from the level of the allocations themselves to the level of dynamic allocation strategies. The goal of this paper is to give a unified perspective on finding spectrum holes without inducing an unacceptable fear of harmful interference. The subsequent use of these spectrum holes as well as the design/enforcement of the regulations are both outside the scope of this paper.

Cognitive radios have been proposed to be the next generation devices that can dynamically share underutilized spectrum [2], [11], [12]. Spectrum sensing has been identified as one of the key enablers for the success of cognitive radios [6], [13]. There has been a lot of work on designing sensing algorithms for cognitive radio systems. Table I gives a brief sampling of some representative single-user sensing techniques. The techniques given in Table I are by no means exhaustive. The reader is encouraged to look into the references within these references for more. In addition to single-user techniques, cooperative approaches have also been proposed. A brief survey of cooperative sensing approaches is given in Table III. However, spectrum sensing is still very much an active area of research and so in this paper we *do not* aim to find the best possible sensing algorithm for identifying spectrum holes. Instead, the goal here is to understand the key concerns in sensing and how different approaches can be compared to each other.

We start by understanding the basic issues in identifying spectrum holes. To do so, it is easier to concentrate on two extreme cases. First consider primary transmitters like television towers that are always communicating to users in their service area. Some of the area around the primary transmitter can never be used (the red area in Figure 1(a)) while areas further away (the green area in Figure 1(a)) could always be used by secondary users. For bands with such primary users, recovering spectrum holes in *space* is the major concern. Contrast this to a system that transmits intermittently but serves the entire area of interest (see Figure 1(b)). For such a band, recovering spectrum holes in *time* is the major concern.

Traditionally, the time-perspective has dominated the literature. The triad of sensitivity, probability of missed detection (P_{MD}), and probability of false alarm (P_{FA}) have been used to evaluate the performance of sensing algorithms [31]. The first two are connected to the level of protection for the primary users while the last is connected to the performance of the secondary user. Meanwhile, the time required to sense provided a measure of the overhead imposed by the sensing strategy. The tradeoff between these four metrics provided the sensing-layer interface to the overall tradeoff between




Fig. 1. (a) Spectrum holes in space. Area around each transmitter (shaded red) can not be used for secondary transmissions. However the shaded green area can be used all the time. (b) Spectrum holes in time. The secondary user cannot transmit while a primary transmission is on (shaded red). A secondary user can hope to reuse the off times of the primary user (shaded green).

the level of protection/safety offered to the primary user and the secondary system performance, but there is not a one-to-one mapping. Secondary system performance is naturally measured using expected throughput, but this makes sense only in the context of a complete system model. Thus, the design problem can be stated as a cross-layer optimization problem of maximizing the data rate while ensuring that the weighted probability of missed detection (the proxy for primary user safety) is bounded [32], [33].

While the cross-layer optimization approach does allow the comparison of disparate sensing strategies, it does so only in the context of a complete system model. Conceptually, this is disturbing because it tightly couples the internals of sensing spectrum holes to the communication strategy used once the holes have been found. We believe that this indicates that the traditional metrics do not represent the right level of abstraction — to have a unified perspective, we need uniform metrics that can compare sensing algorithms (both single-user and cooperative approaches) at the sensing layer itself. The advantage of this approach is that it gives us the freedom to design sensing algorithms without explicitly worrying about higher-layer considerations.⁴ Moreover, these metrics must also allow us to incorporate modeling uncertainties, which can significantly impact the sensing performance.

The need to incorporate uncertainties can easily be seen in the time-domain. For example, exploiting time-domain spectrum holes in the context of Bluetooth and Wireless LAN coexistence has been considered in [34]. The key to exploiting

⁴This is also desirable from a regulatory perspective. Requiring recertification of a complete system each time anything changed would be a tremendous obstacle to innovation. The main goal of regulation is to preserve safety — and this is largely determined by the operation of the sensing-layer.

<i>Detection algorithm</i>	<i>Description of algorithm</i>	<i>What is modeled?</i>	<i>To what gain?</i>
Energy detection [14]–[16]	Get empirical estimate of energy in a frequency band and compare against a detection threshold.	Average power	Baseline detector for comparison.
FFT for DTV pilot signal [17]–[19]	Partial coherent detection using DTV pilot. Filter around pilot to reduce noise power. Use FFT as partial coherent detector for sinusoids.	Signal contains narrowband pilot tone	Sensing time and robustness
Run-time noise calibrated detection [20]	Noise is calibrated during run-time leading to robustness gains.	Asymmetric use of degrees of freedom	Robustness gains.
Cyclostationary detection [21]–[25]	Spectral correlation function reveals peaks at multiples of the modulation rate/pilot frequency.	Signal is modeled as wide-sense cyclostationary	Robustness gains
Dual FPLL pilot sensing [26]	Use two Digital PLLs which are preset to ± 30 kHz around the pilot. Use time to converge as test statistic.	Signal contains narrowband pilot tone	Simplicity of implementation
Eigenvalue based detection [27], [28]	Utilizes the fact that white noise is uncorrelated across samples/antennas while a bandlimited external signal is correlated	Bandlimited primary signal and secondary radio has multiple receive antennas	Sensing time gains
Event-based detection [29], [30]	The detector tries to detect arrival/departure of signals. This technique can be used for identifying time-domain holes.	Primary user ON/OFF durations are much shorter than the time between secondary user movement	Robustness gains

TABLE I

COMPARISON OF REPRESENTATIVE SINGLE-USER SENSING ALGORITHMS FOR DTV DETECTION. THESE ALGORITHMS USE VARIOUS FACETS OF THE TRANSMITTED SIGNAL TO OBTAIN A BETTER DETECTION SENSITIVITY OVER SIMPLE ENERGY DETECTION.

such opportunities in time is the secondary user’s ability to predict the OFF times of the primary users [35], [36]. While these results have established that dynamic spectrum access has the potential to dramatically increase the amount of spectrum available for use, a drawback is that these approaches depend on the detailed model for the primary user’s transmissions. However, real-world uncertainties make it impossible to model real-world transmissions precisely (see [37] for an example from computer networking) and deviations from the assumed model can severely affect the performance of these algorithms leading to interference with the primary system⁵.

The essence of the discussion above is the need for having unifying sensing metrics that capture the right level of abstraction while allowing the incorporation of the relevant modeling uncertainties. It is not too hard to intuit the form of these metrics for the problem of identifying time-domain holes. To get a unified perspective on spectrum sensing, this paper develops the corresponding metrics for the problem of

recovering spectrum holes in space. This problem is non-trivial and is not well understood in the previous literature. A brief comparison of the time-domain and the spatial-domain is given in Table II. The main contributions of this paper are:

- The issue of uncertainty and its modeling is discussed in detail. In particular, the asymmetric nature of the incentives regarding uncertainty-modeling is considered to be at the heart of the dynamic spectrum-recovery problem rather than being merely an annoying complication.
- An explicit approach is given to quantify the Fear of Harmful Interference (F_{HI}) by maximizing the probability of interference under the worst-case environment consistent with the uncertainty model.
- A unified metric, Weighted Probability of Area Recovered (WPAR), is given to measure overall sensing performance. This allows for a simple analysis that decouples different primary users.
- Cooperative approaches are discussed not just under ideal models, but also with the uncertainty that is the unavoidable companion to freedom.
- In-the-field calibration is introduced as a mechanism to reduce environmental uncertainties that have a wider bandwidth than the primary user. Examples of such uncertainties are interference and shadowing.

The rest of the paper is organized as follows: After Section II formally defines a spectrum hole, Section III discusses the relevant metrics to quantify safety (non-interference) for the primary and the area recovered for the secondary. Sec-

⁵This is analogous to open-loop control in stochastic systems [38], [39]. Systems with open-loop control rely heavily on precise and accurate modeling. In contrast, closed-loop control systems can be much more robust to modeling uncertainties. One possible approach to resolve this uncertainty in the spectrum-sharing context is feedback from the primary system. Such feedback can significantly help in robustly exploiting opportunities in the time domain. Opt-in spectrum markets are an extreme case of explicit feedback from primary users [40], but other forms of implicit feedback are also possible. For example, [41] proposes a spectrum-sharing architecture in which the secondary user eavesdrops on a packetized primary user’s automatic repeat request (ARQ) messages to stay within the interference budget of the primary users.

tion IV illustrates the use of the metrics by considering a single-radio approach to finding spectrum holes and reveals the fundamental limitations of the IEEE 802.22 approach to evaluating detectors [42]. The example of the radiometer is used to connect these metrics to earlier perspectives as well as to show how to incorporate the impact of finite sensing times and uncertainty in the fading model. Section V discusses both the potential gains from cooperative detection strategies and their sensitivity to shadowing-correlation uncertainty. Section VI discusses the use of measurements in nearby bands (eg. satellite bands) to enable assisted detection and points to a way to overcome the uncertainty regarding shadowing correlation. Section VII revisits the lessons of this paper and concludes with pointers to future work. To keep the paper accessible to a general audience, mathematical formalism is kept to a minimum. Precise formulations and detailed proofs of the results in this paper are given in [43].

II. DEFINING A SPECTRUM HOLE IN SPACE

In time the definition of a spectrum hole is easy to understand — it is the period of time that the primary is not transmitting. A spectrum hole in frequency is a little more nuanced. If a secondary user finds a frequency band empty (no primary user present in that band), its transmissions can still interfere with primary receivers operating in adjacent frequency bands (due to imperfect filters and analog front-ends). Hence, a spectrum hole in frequency is technically defined as a frequency band in which a secondary can transmit without interfering with any primary users (across all frequencies). For simplicity, we suppress this subtle distinction in this paper and consider a spectrum hole in frequency to be a contiguous frequency band which is not used locally by any primary user. For further simplicity, we will consider only one such frequency band at a time.

Definition 1: Consider a perfect magical detector that tells us whether it is safe to use a particular secondary system at a given point in space-time or not. Denote the output of this detector (the safe-to-transmit region) by $D^* \subset \mathbb{R}^3$ where two of the dimensions represent space and the third represents time. A *spectrum hole* in space-time is defined as an indicator function $\mathbf{1}_{D^*} : \mathbb{R}^3 \rightarrow \{0, 1\}$ defined as

$$\mathbf{1}_{D^*}(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in D^*, \\ 0 & \text{if } \mathbf{x} \in \mathbb{R}^3 \setminus D^*. \end{cases}$$

For further simplicity, we focus on a frequency band which is licensed to a single primary service. The primary transmitters dealing with this particular band are assumed to be distributed over a large geographic area with non-overlapping service areas. For example, consider television bands where primary transmitters⁶ are stationary and have long-lived transmissions. A television station’s transmitter is mounted on a high tower (≈ 500 m) and serves a large radius (≈ 50 km). Further away, the signal from the tower is very weak and a secondary user at such a location can transmit

⁶For simplicity, we ignore the issue of peaceful coexistence with wireless microphones operating in the television band. Such smaller scale primary users introduce additional challenges [44].

without causing interference. Our attention will mostly be focused on a single one of those towers and the area around it.

Figure 2 shows a primary transmitter and a single primary receiver. In the absence of interference, a receiver within the blue circle (Figure 2a) with radius r_{dec} would be able to decode a signal from the transmitter, while a receiver outside the circle would not. To tolerate any secondary users, the primary receiver needs to accept some additional interference. The green circle represents the protected radius (denoted r_p) where decodability is guaranteed to primary receivers. Primary receivers between the two circles may not be able to get service once secondary systems come on, but this is considered to be an acceptable loss of primary user QoS.⁷ Call these “sacrificial zones.” The time-dimension equivalent of $r_{dec} - r_p$ is the short sacrificial time-segment at the beginning of a primary transmission during which secondary users are permitted to cause interference.⁸

Around each protected primary receiver, a no-talk region exists where a secondary user cannot safely transmit. However, this depends on the nature of the secondary transmission. If it has low transmit power, Figure 2a illustrates how the no-talk zones around each receiver can be small. If it has high transmit power, Figure 2b illustrates how the radius of the no-talk zones become much larger. There are two ways to interpret this effect. One approach is to consider the transmit power of the secondary user as its footprint and think of the secondary user as a finite-sized ball (of radius $(r_n - r_p)$). In this approach, the question becomes whether the ball fits into the hole. For simplicity, a second approach is followed here: the secondary user is considered to be a point and the spectrum hole itself is not considered to include those points at which a secondary user would not safely fit.⁹

The overall no-talk area is thus the union of the no-talk regions of all primary receivers. The spectrum hole is the complement of this union. To recover this area, the secondary system must know the locations of all primary receivers (see Figure 3(a)). Since a primary user may know this information, such complete area recovery might be possible with explicit primary participation. In addition, secondary users themselves may be able to determine the locations of receivers for particular TV channels by sensing the TV receivers themselves [45].

However, just because a secondary transmitter *can* safely transmit in a particular location on a particular band does not imply that it should want to do so. After all, close to

⁷This can be viewed as either the loss of service to certain customers of the primary system or as an additional cost of transmit power that must be spent by the primary user to maintain service to all the same customers.

⁸Like its spatial equivalent, this can be viewed as either a loss of QoS for the primary user in the sense of a dropped frame or as requiring the primary user to lengthen its synchronization preamble before commencing data transmission. Without this provision, a secondary user could never transmit due to the fear of primary user reappearance during the secondary transmission.

⁹For simplicity, this discussion assumes a single simultaneous secondary transmission. In practice, the secondary system is likely to contain many transmitters operating simultaneously over a distributed area. Such systems can have their user footprints considered in terms of their power density as shown in [7], [8]. However, the analysis in [44] shows that the first interpretation becomes problematic when we really try to scale to secondary users with very different footprints.




Fig. 2. Weaker secondary users can transmit closer to the protected primary receivers, whereas louder secondary users can only transmit far from the protected primary receivers.

a functioning primary receiver there will usually be a lot of interference from the primary signal itself. It has been proposed that the secondary transmitter may be able to decode the TV signal and use dirty-paper-coding techniques (DPC) and simultaneously boost the primary signal in the direction of interference [46], [47]. However, it has also been shown that this approach is not robust since simple phase uncertainty can significantly lower the performance of such schemes [48]. Other forms of partial information like knowledge of the primary user's codebook are also not useful unless the secondary receiver can actually decode the primary signal and use multiuser detection. Otherwise, it has been shown that the secondary system is forced to treat the primary transmission as noise [49]. Since even marginally decodable primary signals tend to be far louder than the background noise, this suggests that knowledge of the locations of the primary receivers is not that useful in practice.

Consequently, this paper focuses on recovering the region outside the global no-talk zone (r_n) as shown in Figure 3(b). This is the intersection of the spectrum holes corresponding to all possible locations for protected primary receivers. In this picture, knowledge of the relative positions of the primary transmitters and the potential secondary user is key.

III. METRICS AND MODELS

The main task of the secondary system is to determine its relative position with respect to the primary transmitters and to start transmission only if it is reasonably sure that it will not interfere with any of the potential primary receivers. An ideal solution is to require the primary user to register all of its transmitters' positions and for the secondary system to possess the ability to calculate its own position as well as communicate with the registry that records primary user positions.

While the above works for purely spatial spectrum holes, it does not scale well to spectrum holes that span both space and time. It also involves a lot of overhead. Therefore, we must

consider different approaches to detecting spectrum holes and have metrics that can be used to compare their performance.

A. Signal to Noise Ratio (SNR) as a proxy for distance

A natural approach is for the secondary user to estimate the strength of the primary signal as a proxy for the distance from the primary transmitter. The problem then becomes: at what level must the secondary user detect the primary system to be reasonably sure that it is outside the no-talk radius? If p_t (in dBm) is the transmit power of the primary user and α is the attenuation exponent¹⁰, then the secondary user can transmit if the received power from the primary user at the secondary user is less than $p_t - 10 \log_{10}(r_n^\alpha)$ i.e.

$$P \begin{matrix} \text{do not use} \\ \geq \\ \text{use} \end{matrix} p_t - 10 \log_{10}(r_n^\alpha), \quad (1)$$

where P (in dBm) is the received primary power at the secondary radio. In general, P is a random variable and its realization can be computed by taking the log of the empirical average of the square of the received primary signal (See Section IV-B).

The above assumes that a system can perfectly determine its relative position given only the received signal strength and can thereby recover all the area beyond the no-talk radius. In reality, the primary signal may experience severe multipath and shadowing which results in a low received power. Seeing a low power signal, the secondary user may decide that it is outside the no-talk radius while in fact it is inside. Hence, a system must somehow budget for such fading. One possible

¹⁰A commonly used propagation model for DTV signals transmitted from TV towers is given in [50]. The pathloss function described by this model (see Figure 1 in [51]) can be approximated by a continuous piecewise polynomial function. Explicitly, for all the figures in the paper we use an exponent of $\alpha = 3$ for distances below 1 km, an exponent of $\alpha = 2.7$ till 30 km, an exponent of $\alpha = 7.65$ till 100 km, and an exponent of $\alpha = 8.38$ from there on. However, to keep the expressions in the text simple, we use a single polynomial with exponent α for the pathloss function.




Fig. 3. (a) Area within the protected region can be recovered if the positions of the primary receivers can be determined. (b) Global no-talk area defined assuming the primary receivers can be anywhere within the protected region

approach is to introduce a design parameter, Δ (in dB), which is the combined budget for possible fading and shadowing losses. Then, the rule in (1) becomes:

$$P \begin{cases} \text{do not use} \\ \geq \\ \text{use} \end{cases} p_t - (10 \log_{10} r_n^\alpha + \Delta) \quad (2)$$

In (2), the parameter Δ is a constant serving the role of a safety factor. Its value is determined by the desired operating point of the system, and it is fixed at design time. The value of Δ impacts the secondary user's ability to guarantee non-interference to the primary user as well as to recover area for its own operation. If Δ is large then the secondary user acts conservatively and only declares a point usable when the primary signal there is very weak. In normal circumstances such weak signals occur very far from the TV transmitter and the secondary user must forfeit a lot of the area around the primary transmitter (see Figure 4) but it is able to ensure non-interference to the primary user. If Δ is small, there is a chance that the secondary user will not even sense moderately faded primary signals. The secondary user will then be interfering with the primary user more often but will forfeit a smaller area (see Figure 4). This tradeoff needs to be captured in the appropriate metrics.

B. Traditional sensing metrics

We briefly review the traditional triad of sensing metrics (sensitivity, P_{FA} , and P_{MD}) and motivate the need for system-level metrics for the problem of identifying spatial spectrum holes.

Any sensing algorithm can be thought of as a system (black box) with inputs, outputs and control knobs. The input to the system is the received signal, and the output is the decision whether the band is usable or not. The control knobs are design parameters like detector threshold, sensing time, etc. Traditionally, the performance of such a system is characterized by its *Receiver Operating Characteristic* (ROC)




Fig. 4. If the budget for multipath and shadowing is small (Δ small), then the secondary user does not forfeit much area beyond the true no-talk zone. If the budget for multipath and shadowing is large (Δ large), then the secondary user forfeits a lot of area outside the no-talk zone.

curve. The ROC of a detector is the curve that plots the P_{MD} as a function of the P_{FA} for a fixed sensing time, and fixed operating SNR [52]. An alternate performance metric for a detector is its *sensitivity*. The sensitivity of a detector is the lowest value of the operating SNR for which the detector satisfies a given target P_{FA} and P_{MD} .

The overhead for a detector is traditionally measured by the sensing time required to achieve a target P_{FA} , P_{MD} at a given SNR . This is called the *sample complexity* of the detector. The sample complexity and sensitivity are tightly coupled — if we want to improve the sensitivity of the detector, we must increase the sample complexity and hence incur a larger sensing overhead.

An important functional requirement for detectors operating at low $SNRs$ is *robustness* to uncertainties in the system. Un-

certainties can be broadly divided into two classes — device-level uncertainties (like uncertainty in the noise power) and system-level uncertainties (like uncertainty in the shadowing distribution). It was shown in [16] that the traditional metrics can be suitably modified to characterize detector robustness to device-level uncertainties. This was done by considering worst case P_{FA} , P_{MD} over the set of uncertain distributions. Furthermore, it was shown that detectors have fundamental SNR thresholds called SNR walls below which detection is impossible even if the sensing time is increased to infinity. This showed that under device-level uncertainties, we must consider both sensitivity and the detector’s SNR wall as a measure of performance.

Now, the remaining question is: how do we deal with system-level uncertainties? The dominant current approach to deal with system-level uncertainties like uncertainty in shadowing is to incorporate them into the specifications for the system. For instance, to account for possible deep fades, the 802.22 working group specifications require detectors to have a sensitivity of -116 dBm (-20 dB SNR) [42]. This corresponds to a safety margin of roughly $\Delta = 20$ dB [14].

There are two fundamental problems with this approach. First, this approach is very conservative and leads to severe overheads (20 dB \approx 110 km). In most situations detectors do not face such severe fading and hence they are forced to not use the band even though they are well outside the no-talk radius. Secondly, this approach of specifying a sensitivity requirement is not compatible with cooperative sensing approaches. It is clearly hard to even define what sensitivity means for a whole group of radios [53]. What if one of them is faded and the other is not?

C. New system-level metrics

In the previous section we showed why the traditional sensing metrics fail to capture the right level of abstraction between the sensing and communication. Table II lists the quantities/modeling philosophy that we want to capture with appropriate metrics. For the problem of recovering time-domain holes these quantities are well understood (listed in the second column of Table II). The analogous quantities in the spatial domain are listed in the third column of Table II.

We now give two new system-level metrics — safety to the primary user and sensing overhead given by the loss in available area. The metrics have been defined to capture the essence of the discussion in Table II.

1) *Safety*: The first idea for a safety metric is to just calculate the probability of interference. However, this is a metric that is open to serious abuse. A secondary system might do no sensing and just assume that its users will be uniformly placed in a large area (much larger than the footprint of the TV station). Hence the probability of a user landing within the no-talk area is very small and the secondary system can claim compliance with a low target probability of interference.

Such a metric for safety is essentially no better than the secondary system telling the primary user “trust me, my users are not going to be close enough to interfere with you.” The primary user has no reason to trust the *a priori* user-deployment model of the secondary system once the secondary

products are in the marketplace. There is an asymmetry here: the secondary operator might very well have a uniform-area business model in mind, but the primary user fears that the secondary operator will end up deploying the system close to the primary’s receivers since that is where the people are. A metric that accurately captures the primary’s fear of harmful interference must somehow assume the worst-case deployment of secondary users.

Similarly, there is no reason to completely trust the fading model. A detector could end up operating in line-of-sight environments or it could be deeply shadowed. For example, the secondary operator may propose roof-top static installations (with very little shadowing) of its access devices thinking that people will be using it to get Internet access in single-family homes. However, people living in apartment buildings might also start buying the devices. Some users might notice that system performance improves if they bring their devices indoors (becoming shadowed from primary transmissions). A new multiplayer video game might even arise that encourages people to use the device inside their minivans while driving around town. The primary user will not trust the secondary operator to alienate its own paying customers and it is hard to perfectly anticipate the environment of the future.

The following definition captures these model uncertainties.

Definition 2: Assume that the secondary user runs a spectrum-sensing algorithm that outputs a binary decision D about the state of the primary band: 1-used/0-unused. The *probability of potential interference* $\mathcal{P}_{F_r}(D = 0|r_{actual} = r)$ at radius $r \leq r_n$ is the probability that a secondary user is within the no-talk region and declares¹¹ that the band is “unused”. Here F_r is the probability distribution of the combined multipath and shadowing-induced fading at a distance r from the primary transmitter.

The exact value of this probability depends on the assumed model for shadowing and multipath. The primary users (and regulators) only trust that the true distribution is within the set \mathbb{F}_r . Hence the *Fear of Harmful Interference* (F_{HI}) is defined as:

$$F_{HI} = \sup_{0 \leq r \leq r_n} \sup_{F_r \in \mathbb{F}_r} \mathcal{P}_{F_r}(D = 0|r_{actual} = r). \quad (3)$$

The outer supremum reflects the uncertainty in secondary user deployments and the inner supremum reflects the uncertainty in the distribution of the fading. Explicit models for these uncertain distributions are discussed in Section III-D.

There is an analogous safety metric for spectrum holes in time where the goal is to reuse the primary user’s OFF time while avoiding harmful interference in ON times. In addition to the fading uncertainty, the distribution of the inter-transmission times of the primary transmitters must also be viewed as uncertain (see e.g. [37]) to preserve the freedom of action of the primary system’s users. In addition, the relative starting time of the potential secondary transmissions is also viewed as uncertain just as the secondary position in space is considered uncertain.

¹¹This does not necessarily mean that a secondary radio will actually transmit and cause interference.

<i>Quantity of interest</i>	<i>Time domain</i>	<i>Spatial domain</i>
Interference margin	Permissible duration of secondary interference at the start of primary user's ON period	Marginal area relinquished by primary users to allow secondary operation
Modeling uncertainty	Distributional uncertainty in the primary users' ON/OFF periods	Distributional uncertainty in the primary signal's fading/shadowing
Scenario for computing safety metric	Worst-case overlap between primary's and secondary's transmissions	Worst-case placement of secondary users within the no-talk region
Performance metric	Fraction of primary user's OFF period recovered for secondary transmission	Area outside the primary's no-talk region recovered for secondary transmissions
Overhead	Sensing time	Area outside the no-talk region that cannot be recovered

TABLE II
CORRESPONDENCES BETWEEN THE QUANTITIES OF INTEREST IN THE TIME AND SPATIAL DOMAINS.

2) *Performance*: Next we consider a metric to deal with the secondary user's performance — its ability to identify spectrum opportunities. If there were only a single primary transmitter, every point at a radial distance $r > r_n$ would be a spectrum opportunity. For any detection algorithm, there is a probability associated with identifying such an opportunity, called the probability of finding the hole P_{FH} :

$$P_{FH}(r) = \mathcal{P}_{F_r}(D = 0 | r_{actual} = r), \quad r > r_n. \quad (4)$$

In reality, secondary users might also be uncertain about the shadowing and fading distributions. In this case the secondary users can compute performance assuming the worst-case distribution in their uncertainty set. This uncertainty set is typically much smaller than the uncertainty set used in (3) to compute the safety performance to the primary user. This is because the primary user does not trust the secondary users' deployment model and hence assumes a larger uncertainty set. On the other hand the secondary users know their deployment model accurately as there is no incentive for the secondary users to lie to themselves. So, they can work with a much smaller uncertainty set to compute performance. For simplicity, we just shrink the uncertainty set to a single point and assume complete knowledge of the combined shadowing and fading distribution, F_r .

The goal is to combine the probabilities $P_{FH}(r)$ into a single performance metric that allows a comparison among different sensing algorithms. One choice is the underlying utility of the secondary system, like the total throughput or profit. However, such holistic utility functions are intertwined with the system architecture and business models along with assumptions regarding the placement of all the primary transmitters and the population distribution of potential customers. It is useful to find an approximate utility function that decouples the evaluation of the sensing approach from all of these other concerns.

We make the reasonable assumption that secondary utility will increase whenever additional area is recovered by the sensing algorithm. Since we would like to decouple the sensing metric from the detailed model for primary deployments, it is useful to be able to state it in terms of a single primary transmitter. The difficulty is that if there is only a single

primary transmitter, the total area of the spectrum hole is infinite. We propose a discounted-area approach analogous to the present-value of consumer utility proposed by [54].

Definition 3: The *Weighted Probability of Area Recovered* (WPAR) metric is

$$WPAR = \int_{r_n}^{\infty} P_{FH}(r) w(r) r dr, \quad (5)$$

where $w(r)$ is a weighting function that satisfies $\int_{r_n}^{\infty} w(r) r dr = 1$.

The numerical results in this paper have been computed using an exponential weighting function, $w(r) = A \exp(-\kappa r)$. While similar results can be obtained for any other weighting function, the exponential weighting is not unreasonable for the following reasons.

- Since TV towers are often located around areas of high population density, areas around the no-talk region are more valuable in terms of deploying a secondary system than areas far away. This can be viewed as a spatial analogy to “banker's discounting” in which money in the future is worth progressively less in present units. By Sutton's law¹², the economic value of an area is proportional to the number of potential customers there. Population densities are often modeled as decaying exponentially as one moves away from the central business district [56].
- As we move away from any specific tower, there is a chance that we may enter the no-talk zone for another primary tower transmitting on the same frequency. This can be viewed as a spatial analogy to “drug-dealer's discounting” in which money in the future is worth less than money in the present because it is uncertain whether the drug dealer will survive into the future because of the arrival of the police or a rival gang [57].

Figure 5 shows the locations of TV transmitters for Channel 30 all around the United States [58]. In keeping with the current rural deployment assumptions of IEEE 802.22, we just consider “drug-dealer's discounting” here and this sets

¹²When asked why he robbed banks, the famous bank robber Willie Sutton is believed to have said “because that is where the money is” and so this general principle has been named after him [55].

the value of $\kappa = 2 \times 10^{-5} m^{-1}$ for the paper, given the other parameters that are commonly used for digital television signals: primary transmit power $p_t = 90$ dBm, no-talk radius¹³ $r_n = 150.3$ km, and a piecewise polynomial propagation model fitted to match Figure 1 in [51].

When dealing with intermittent primary users (i.e. trying to recover holes in time), the goal is to reuse the OFF time while minimizing the sensing time. To understand the relative burden of the sensing time, we need to appropriately weigh recovered opportunities in time. ‘‘Drug-dealer’s discounting’’ is appropriate since potential opportunities in the future may never materialize because there is a chance of the primary user re-appearing before then. Thus, there is a Weighted Probability of Time Recovered (WPTR) metric that is analogous to the WPAR metric proposed in this section. In the interests of space, this metric is not pursued further here.

D. Models for fading uncertainty

The received primary signal strength P (in dBm) can be modeled as $P = P_t - (l(r) + S + M)$, where P_t is the power of the transmitted signal, $l(r)$ is the loss in power due to attenuation at a distance r from the primary transmitter, S is the loss due to shadowing and M is the loss due to multipath fading. Unless specifically mentioned, we assume that all powers are measured in dB scale. We assume that $l(r) = 10 \log_{10}(r^\alpha)$, and α is the true attenuation exponent.¹⁴

1) *Nominal model*: For convenience, S and M are assumed to be independent of r and to follow a nominal model for $S + M$ that is Gaussian ($S + M \sim \mathcal{N}(\mu_S, \sigma^2)$) on a dB scale. This implies that $P \sim \mathcal{N}(\mu(r), \sigma^2)$, where $\mu(r) = P_t - (l(r) + \mu_S)$. This is the distribution used to compute the WPAR as in (5). For the plots in this paper, $\mu_S = 0$ dB and the standard deviation $\sigma = 5.5$ dB were chosen to match standard assumptions in the IEEE 802.22 literature [51].

2) *Quantile models*: To compute F_{HI} , we cannot always use the nominal model for shadowing and multipath as it is important to model the fact that the primary user does not trust this model completely. Instead, it is possible that the primary user trusts only a quantized version (or a coarse histogram) of the fading distribution. Mathematically, we model this as a class of distributions (\mathbb{F}_r) that satisfy given quantile constraints.

Definition 4: A single quantile model \mathbb{F}_r is a set of distributions for the received signal power defined by a single number $0 \leq \beta \leq 1$ and a function of r denoted $\gamma(r, \beta)$. A distribution $F_r \in \mathbb{F}_r$ iff

$$\mathcal{P}_{F_r}(P < \gamma(r, \beta)) = \beta. \quad (6)$$

A k -quantile model is a set of distributions \mathbb{F}_r for the received signal power defined by a list of numbers ($\beta_1 < \beta_2 < \dots < \beta_k$) and a corresponding list of functions

$(\gamma_1(r, \beta_1), \dots, \gamma_k(r, \beta_k))$. A distribution $F_r \in \mathbb{F}_r$ iff $\forall i \leq k$

$$\mathcal{P}_{F_r}(P < \gamma_i(r, \beta_i)) = \beta_i. \quad (7)$$

For consistency, the quantiles are chosen so that the nominal Gaussian $\mathcal{N}(\mu(r), \sigma^2)$ is always one of the possible distributions for P .

$$\gamma(r, \beta) = \mathcal{Q}^{-1}(1 - \beta)\sigma + \mu(r), \quad (8)$$

where $\mathcal{Q}^{-1}(\cdot)$ is the inverse of the standard Gaussian tail probability function.

Figure 6 shows a picture of the distributions allowed under the quantile model (5 learned quantiles) defined in this section. The set of allowed Cumulative Distribution Functions (CDF’s) for P under our quantile model is precisely the set of all possible non-decreasing curves sandwiched between the upper and lower bounds shown in Figure 6. The dashed (black) curve in the figure shows the nominal Gaussian CDF for P , and the quantile constraints can be thought of as samples of the nominal CDF (the triangle points (in red) in the figure).




Fig. 6. The quantile model for the received signal power (P) distribution. The dashed (black) curve is the nominal Gaussian CDF for P , and the triangle points (in red) show the quantile constraints on the CDF. The dashed-dotted (magenta) curve is the upper bound and the solid (blue) curve is the lower bound on the allowable CDF for P . The actual CDF can lie anywhere in between, and must pass through the 5 triangle points (quantile constraints).

IV. SINGLE-RADIO SENSING PERFORMANCE

The tradeoff between F_{HI} and $WPAR$ depends on the detector used by the secondary user. We start with a hypothetical detector that meets the current specification specified in the IEEE 802.22 process. The issue of finite sensing time is illustrated next through the example of a radiometer. A similar analysis could be carried out for any other detection algorithm and so the role of uncertain fading distributions is investigated using an ideal detector with an infinite sensing time.

¹³This corresponds to WRAN basestations in 802.22. Using 36 dBm for secondary transmitters gives the 150.3 km radius [51].

¹⁴We could include an uncertainty model for the attenuation exponent since the antenna heights can vary and include this in the computation of F_{HI} . However, for simplicity we assume complete knowledge of the attenuation exponent in this paper.



Fig. 5. Location of transmitters for Channel 30 (566-572MHz) plotted using Google Maps.

A. Evaluating an ideal -116dBm detector

The currently understood detector specifications in the IEEE 802.22 working group require any proposed sensing algorithm to be able to detect digital television signals at -116 dBm to a probability of mis-detection $P_{MD} = 0.1$ and probability of false alarm $P_{FA} = 0.1$ [42].¹⁵ We now show that detectors based on such specifications lead to very poor area recovery and also do not guarantee safety beyond the 0.1 level without additional unspoken assumptions.

Suppose a detection algorithm meets the -116 dBm, $P_{MD} = 0.1$, $P_{FA} = 0.1$ specification. Since only the -116 dBm level is specified, it is natural to assume that the primary user only has confidence in a single quantile that corresponds to that level, i.e., $\mathcal{P}(P < -116) = \beta(r)$, where $\beta(r) = Q\left(\frac{-116 - \mu(r)}{\sigma}\right)$.

¹⁵The specified sensitivity of -116 dBm was based on the observation that it is easier (think shorter verification times) to verify a probability specification of 0.9 than it is to verify a probability specification of 0.999. Furthermore, there was an expectation that detector performance would monotonically increase with increased received power. Hence, a detector that demonstrated a probability of detection of 0.9 at -116 dBm would (hopefully) demonstrate a much higher detection probability at -110 dBm [59].

Let \mathcal{D} denote the set of all detection algorithms satisfying the IEEE 802.22 specifications. Then, the fear of harmful interference is,

$$\begin{aligned}
 F_{HI} &= \sup_{0 \leq r \leq r_n} \sup_{F_r \in \mathbb{F}_r} \sup_{D \in \mathcal{D}} \mathbb{E}_{F_r} [P_{MD}^D(P)] \\
 &\stackrel{(a)}{=} \sup_{0 \leq r \leq r_n} \sup_{D \in \mathcal{D}} P_{MD}^D(-\infty) \mathcal{P}(P < -116) + \\
 &\quad P_{MD}^D(-116) \mathcal{P}(P \geq -116) \\
 &\stackrel{(b)}{=} \sup_{0 \leq r \leq r_n} \sup_{D \in \mathcal{D}} [(1 - P_{FA}^D) \mathcal{P}(P < -116) + \\
 &\quad P_{MD}^D(-116) \mathcal{P}(P \geq -116)] \\
 &\stackrel{(c)}{=} \sup_{0 \leq r \leq r_n} [\beta(r) + 0.1(1 - \beta(r))] \\
 &= \sup_{0 \leq r \leq r_n} (0.9\beta(r) + 0.1) \\
 &\stackrel{(d)}{=} 0.9\beta(r_n) + 0.1.
 \end{aligned}$$

In the above chain of equalities the superscript D is used to denote a detection algorithm from the class of allowed

detection algorithms \mathcal{D} . Equality (a) follows from the fact that the maximizing distribution $F_r^* \in \mathbb{F}_r$ corresponds to placing a mass of $\beta(r)$ at $-\infty$ and $(1 - \beta(r))$ at -116 dBm. This is the maximizing distribution irrespective of the actual detection algorithm $D \in \mathcal{D}$. This is because $P_{MD}^D(p)$ is a monotonically decreasing function of p , for all $D \in \mathcal{D}$. Equality (b) follows from the fact that $P_{MD}^D(-\infty)$ is the mis-detection probability when the signal is absent ($p = -\infty$). This corresponds to the event when noise-only received signal samples do not cause a false-alarm. Hence, $P_{MD}^D(-\infty) = 1 - P_{FA}^D$. Equality (c) follows from the fact that $\sup_{D \in \mathcal{D}} (1 - P_{FA}^D) = 1$, and $\sup_{D \in \mathcal{D}} P_{MD}^D(-116) = 0.1$. Finally, equality (d) follows from the fact that $\beta(r)$ is a monotonically increasing function.

Now, for any $D \in \mathcal{D}$, the probability of finding a hole is given by

$$\begin{aligned} P_{FH}(r) &= \mathbb{E}_P [\mathcal{P}(T^D(\mathbf{Y}) < \lambda | P)] \\ &= \mathbb{E}_P [P_{MD}^D(P)] \\ &\stackrel{(e)}{\leq} P_{MD}^D(-\infty) \mathcal{P}(P < -116) + \\ &\quad P_{MD}^D(-116) \mathcal{P}(P \geq -116) \end{aligned} \quad (9)$$

$$\stackrel{(f)}{\leq} \beta(r) + 0.1, \quad \text{for } r > r_n. \quad (10)$$

The bound in (e) follows from the fact that the function $P_{MD}^D(p) \leq P_{MD}^D(-\infty)$, for $-\infty < p < -116$, and $P_{MD}^D(p) \leq P_{MD}^D(-116)$, for $-116 \leq p < \infty$. These inequalities follow from the fact that $P_{MD}^D(p)$ is a monotonically decreasing function of p . The bound in (f) follows from observing that $P_{MD}^D(-\infty) \leq 1$, $\mathcal{P}(P < -116) := \beta(r)$, $P_{MD}^D(-116) \leq 0.1 \quad \forall D \in \mathcal{D}$, and $\mathcal{P}(P \geq -116) \leq 1$.

Shockingly, there is a benefit from missed detections above! This suggests that a clever detector designer would do well to introduce intentional missed detections to improve performance while still meeting the official specification. This calls into question the unspoken assumption that deployed detector implementations would have better probabilities of missed detection when the primary signal is stronger than -116 dBm.

Using (10) in the definition of WPAR (See (5)) and applying our nominal model gives a $WPAR \leq 0.16$. This clearly shows that while the -116 dBm requirement seems very conservative, the detector specification is actually not very safe and simultaneously has a poor area recovery irrespective of the actual detector used. In the worst-case, the signal can indeed fall as low as -116 dBm at the no-talk radius. However in the average case the signal is a lot stronger and this leads to a lot of valuable area going unrecovered.

B. The radiometer

In the rest of the paper we assume that the received signal is sampled and hence we work in discrete time for simplicity. The radiometer collects the samples of the received signal $Y[n]$, computes its empirical power and compares it to a detection threshold. The test-statistic for the radiometer can be written as

$$T(\mathbf{Y}) = \frac{1}{N} \sum_{n=1}^N |Y[n]|^2 \underset{D=0}{\overset{D=1}{\geq}} \lambda, \quad (11)$$

where λ is the design parameter called the detector threshold. Here, $Y[n] = X[n] + W[n]$, where $X[n]$ is the faded primary signal at time n and $W[n]$ is the background noise at time n . For convenience assume that all $W[n]$ are independent and identically distributed as $\mathcal{N}(0, \sigma_w^2)$. Also, let N be the total number of samples that are collected for sensing.

The average power of the received primary signal is given by $P = 10 \log_{10} \left[\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N |X[n]|^2 \right]$ (in dBm). The WPAR does not depend on any uncertainty and so

$$P_{FH}(r) = E[\mathcal{P}(T(\mathbf{Y}) < \lambda | P = p)], \quad (12)$$

where the outer expectation is with respect to the nominal Gaussian distribution $P \sim \mathcal{N}(\mu(r), \sigma^2)$. Substituting (12) in (5) gives the WPAR for the radiometer.

The analysis to compute the fear of harmful interference for the radiometer under the single quantile uncertainty model is similar to the analysis in Section IV-A, and hence is not repeated here. Note that the secondary user has two parameters to adjust. It can adjust the threshold λ on its own and it can negotiate with the regulator/primary user regarding the appropriate value for β . We assume that it does both and chooses the optimal β and λ so as to maximize the WPAR. This can be done numerically. Figure 7 shows the resulting safety/performance tradeoff for a single radio with both a finite and infinite number of samples. For the case when $N = \infty$, it is easy to see that the optimal choice for $\lambda = \gamma(r_n, \beta) = \sigma \mathcal{Q}^{-1}(1 - \beta) + \mu(r_n)$. From Figure 7 we can see that the impact of the uncertainty is substantial when the sensing time is finite.




Fig. 7. Performance of a perfect detector (infinite samples) as compared with a radiometer using a finite number of samples.

C. The value of additional samples

Next we look at the gap between the safety-constrained performance with only a single (but optimized) trusted quantile and what can be achieved with the entire fading distribution

being trusted. The first step is to increase the sensing duration. Figure 8 shows the performance as the number of samples N is scaled but the F_{HI} of the system is constrained to be 10^{-3} . An infinite number of samples leads to a perfect detector, and it turns out that having a single trusted quantile leads to the same performance as having complete distributional knowledge. This is because that single quantile can be chosen in an optimal fashion based on the target F_{HI} itself. Hence the two curves achieve the same WPAR value as the number of samples are scaled up. However, they need different numbers of samples. If the entire distribution were trusted, a single radio only needs $\sim 10^4$ samples whereas $> 10^7$ samples are needed if only a single quantile can be trusted.




Fig. 8. Performance of a radiometer with finite samples approaches the performance of the perfect detector as the number of samples is increased.

D. The value of additional consensus

The previous section shows that there is a clear value to agreeing on a single model for the entire distribution. However, this is likely to be impossible in practice. Instead, suppose that the primary user, secondary user, and regulators agreed on a few quantiles of the fading distribution instead of a single one.

Figure 9 shows the fear of harmful interference (F_{HI}) for a fixed WPAR as the number of quantiles is increased while the sensing time is kept constant. Two methods for quantile selection are compared. In the first method the quantiles are chosen uniformly (e.g. if three quantiles were needed, select the 1/4th, 1/2 and 3/4th quantile). In the second method, the best additional quantile is chosen greedily given the choice of the previous quantiles. Both methods approach the same limit, but the greedy choice clearly performs better. The threshold used in this plot corresponds to a F_{HI} of 0.1 if the entire distribution is trusted. A moderate number of quantiles (~ 10) are needed for the safety to be reasonably close to the safety with complete distributional knowledge, for the same detection threshold.

If the F_{HI} were to be held constant, the WPAR performance would improve instead. By gaining additional consensus regarding the fading distribution, the sensing threshold can be set more aggressively without increasing the fear of harmful interference. This aggressive threshold in turn increases the WPAR.




Fig. 9. The F_{HI} of an energy detector with 100 samples but distributional uncertainty approaches the F_{HI} with no uncertainty as the number of known quantiles is increased. Different ways of choosing the quantiles show different performance. The threshold used in this plot corresponds to an F_{HI} of 0.1 without distributional uncertainty.

E. The value of improved detection algorithms

From Figure 7 we can see that even a perfect radiometer recovers only a 0.37 fraction of the weighted area for a safely low F_{HI} ($\approx 10^{-2}$). It is tempting to believe that performance could be improved by considering more powerful detectors like pilot detectors and cyclostationary feature detectors [60]. It has certainly been shown that pilot detectors and cyclostationary feature detectors are more robust to uncertainties in the noise process [16]. However, at best such single-user detectors can achieve the performance of a perfect radiometer, but this is limited due to the need to budget for deep fades.

V. COOPERATION

One possible approach to solve the problem mentioned in Section IV-E is to use the sensing results from multiple nearby radios to make a decision on whether the band is free to use or not. This mirrors previous research in cooperative communications and sensor networks [61], [62]. Several groups have proposed cooperation among cognitive radios as a tool to improve performance [13], [63]–[69]. Table III lists the major research themes in the area of cooperative spectrum sensing and representative references. Gains from cooperation can either be viewed as diversity gains where multiple radios reduce the collective probability of getting a bad fade [13], [63] or as a mechanism to reduce sensing overhead [64], [65],

[69]. Dealing with uncertainty (in the form of correlated data measurements and/or malfunctioning/malicious radios) forms a major component of this research. In addition, the design of optimal cooperative sensing schemes under various constraints (communication/synchronization constraints for example) is also an active area of research.

We believe that the most significant gains from cooperation (from the standpoint of recovering spatial holes) are diversity gains. Hence we look at cooperation as a tool to increase WPAR. We assume that a group of M cognitive radios are listening to the primary signal on a given frequency band. This group makes a common decision on whether the band is free to use or not. For simplicity, we assume that each radio gets a perfect estimate¹⁶ of the received primary power P_i (in dB) $i = 1, \dots, M$. We make this assumption to isolate the gains due to cooperation from those due to a longer effective sampling time.

Each of the received signal strengths is written as $P_i = p_t - (10 \log_{10} r_i^\alpha + S_i + M_i)$, where p_t is the transmit power of the primary signal, r_i is the distance from the i th radio to the TV tower, and S_i and M_i are respectively the losses due to shadow and multipath fading at the i th radio. All cooperating radios are assumed to be located at approximately the same distance from the TV tower, i.e., $r_i = r$ for all $i = 1, 2, \dots, M$. This models the case when the scale of cooperation is much smaller than the scale of the primary transmissions.¹⁷ This assumption also guarantees that all the cooperating radios are trying to identify the same spectrum hole in space.

To start with, shadowing and multipath are modeled to be independent across the different radios. It is safe to assume that the $\{M_i\}$ are independent¹⁸ of each other since multipath is independent at distances on the order of a few wavelengths [83]. By contrast, shadowing is independent only on a much larger spatial scale [84]. Even though independence might not be an accurate modeling assumption, we first analyze cooperative gains under this best-case assumption. Then, Section V-D computes the loss in performance if the shadowing is not independent.

A. Maximum-likelihood detector: soft-decision combining

Our goal is to find the optimal estimate of the distance r , given the vector of received power observations (P_1, P_2, \dots, P_M) . When the model is completely known, the optimal detector is the ML detector. We assume a nominal Gaussian model for both the shadowing and multipath distribution, i.e., $P_i \sim \mathcal{N}(\mu(r), \sigma^2)$, where $\mu(r)$ is some deterministic monotonically decreasing function of r . Under this model, the mean of the received power is dependent on its distance from the TV tower and the standard deviation is independent of the

¹⁶We can obtain a perfect received primary power estimate by running a radiometer for very long sensing times, i.e., $N \rightarrow \infty$.

¹⁷In reality, the radial footprint of the cooperating radios has to be dealt with as a minor increase in the no-talk radius r_n . However, we assume that the footprint of cooperation is much smaller than the margin $(r_n - r_p)$ and thus ignore this small effect.

¹⁸This is not true strictly speaking. In general, M_i 's are conditionally independent given the shadowing environment since the shadowing environment can determine if there is or is not a strong line-of-sight path. However, we are assuming indoor operation and so there are no line-of-sight paths.

distance from the tower. Using this model, it is easy to see that the ML detector is equivalent to

$$\frac{1}{M} \sum_{i=1}^M P_i \underset{D=0}{\overset{D=1}{\gtrless}} \lambda. \quad (13)$$

This detector computes the average received signal power on a dB scale (This is an example of a soft-decision combining rule since the radios have to send their received power values to a central combiner rather than just sending 1-bit decisions) and compares it to a threshold λ . The frequency band is declared free if the mean signal power is less than λ .

Assuming that $P_i \sim \mathcal{N}(\mu(r), \sigma^2)$, we have $\frac{1}{M} \sum_{i=1}^M P_i \sim \mathcal{N}(\mu(r), \frac{\sigma^2}{M})$. Therefore if we assume that the primary user also trusts the nominal model,

$$\begin{aligned} F_{HI} &= 1 - \mathcal{P} \left(\frac{1}{M} \sum_{i=1}^M P_i \geq \lambda | r_{actual} = r_n \right) \\ &= 1 - \mathcal{Q} \left(\frac{\lambda - \mu(r_n)}{\frac{\sigma}{\sqrt{M}}} \right). \end{aligned} \quad (14)$$

The detector threshold λ must be chosen such that $F_{HI} \leq F_{HI}^{target}$. Hence (14) gives

$$\lambda = \frac{\sigma}{\sqrt{M}} \mathcal{Q}^{-1}(1 - F_{HI}^{target}) + \mu(r_n). \quad (15)$$

For this choice of λ , the probability of finding a hole is

$$\begin{aligned} P_{FH}(r) &= \mathcal{P} \left(\frac{1}{M} \sum_{i=1}^M P_i \leq \lambda | r_{actual} = r \right) \\ &= 1 - \mathcal{Q} \left(\frac{\lambda - \mu(r)}{\frac{\sigma}{\sqrt{M}}} \right). \end{aligned} \quad (16)$$

The WPAR can be computed by substituting (16) into (5). Figure 10 shows the performance of the maximum likelihood detector for several values of the number of cooperating radios M . It is clear that the performance significantly improves even with a few cooperating radios ($M=5$). If $M \rightarrow \infty$, all the area is eventually recovered.

B. Soft-combining with uncertain models

The improvements with cooperation illustrated in Figure 10 assume complete consensus regarding the fading distribution. In reality it is likely that the primary user of the channel does not trust the nominal Gaussian models for shadowing and fading distributions. The cost of addressing this distrust of primary users is a reduced performance for the same value of safety. For now, the independence assumption for fading across different users is maintained.

Under the independent fading assumption, it is illustrative to use the quantile models discussed in Section III-D for each received power P_i . Start with a single quantile that can be optimized. Let the class of marginal distributions satisfying the β th quantile constraint be denoted by \mathbb{F}_r . For simplicity, consider the two cooperating radios case, i.e., $M = 2$.

The maximum-likelihood estimate detector under uncertain fading distributions (even for a single-quantile uncertainty model) does not even make sense. Hence, we do not attempt to

Research Theme	Main idea/goals	References
Cooperation as diversity	Cooperation can be seen as providing diversity gains by reducing sensitivity requirements for individual radios.	[13], [63]
Cooperation as gains in degrees of freedom	Cooperation can be seen as reducing sensing time or lowering false alarms for the same level of detection.	[64], [65], [69]
Impact of/Dealing with correlation	Determining the impact of channel correlation on cooperation gains as well as mechanisms of dealing with correlation uncertainty.	[63], [68], [70]
Impact of/Dealing with malicious/lying users	Determining the impact of incorrect sensing responses and mechanisms for weeding out misbehaving users.	[63], [71]–[73]
Cooperation and Communication	Determine the impact of communications/synchronization constraints on cooperation performance.	[64], [65], [74]–[76]
Fusion rules	Investigation of various soft/hard combining rules.	[67], [69], [77]–[80]
Utilizing sparsity/multiband information	Utilize multiple frequency bands for cooperative gains.	[81], [82]

TABLE III
DESCRIPTION OF VARIOUS RESEARCH THRUSTS IN THE AREA OF COOPERATIVE SPECTRUM SENSING.




Fig. 10. Performance of the ML detector in (13) with complete knowledge of the fading/shadowing distribution.

solve for the best possible detector under modeling uncertainties. Instead, we continue to work with the averaging detector given in (13). As discussed in Section V-A, this detector is the ML detector under perfectly modeled Gaussian fading.

For this detector, we can show that for a given choice of quantile β , the best choice of λ that minimizes F_{HI} is $\lambda = \gamma(r_n, \beta)$, where $\gamma(r_n, \beta)$ is the β th quantile threshold in (8). For this choice of λ , the fear of harmful interference is given by

$$\begin{aligned}
 F_{HI} &= \sup_{0 \leq r \leq r_n} \sup_{F_r \in \mathbb{F}_r} \mathcal{P}_F \left(\frac{P_1 + P_2}{2} \leq \lambda | r_{actual} = r \right) \\
 &= 1 - (1 - \beta)^2.
 \end{aligned} \quad (17)$$

The expression for F_{HI} in (17) can easily be derived




Fig. 11. The averaging detector for two user cooperation under the single-quantile fading model. The solid (black) box in the perimeter represents the P_1, P_2 plane, the solid (blue) line represents the 2-user ML detector in (13), and the dashed (red) lines represent the quantiles describing the distribution of P_1 and P_2 . The shaded area represents the region of the received power pairs (P_1, P_2) for which the detector declares the band unused and the unshaded area represents the region where the detector declares the band as used. The threshold p_{th} in the figure is used to denote $\gamma(r_n, \beta)$.

graphically from Figure 11. In this figure, the P_1, P_2 plane is divided into four quadrants as marked by the dashed-dotted lines (red). The single quantile constraint on the marginal distributions can be written as probability mass constraints within each quadrant. The averaging detector in (13) for a fixed λ can be drawn as a straight line dividing the P_1, P_2 plane into two half planes (the solid (blue) line in Figure 11). If the received power (P_1, P_2) falls in the shaded region, the band is declared ‘free to use’, otherwise the band is declared ‘used’. Hence, the probability of harmful interference is the supremum of the probability mass in the shaded region, where the supremum is taken over all distributions $F \in \mathbb{F}_\beta$. Similarly,

the probability of finding a hole is the probability mass in the unshaded region, under the nominal distribution.

If $\lambda < p_{th}$, where $p_{th} = \gamma(r_n, \beta)$ (the detector line is on the left of the black dot in the figure), then F_{HI} is the sum of probabilities in quadrants *II*, *III*, and *IV*. This is because one can always choose a distribution that satisfies the quantile constraints and puts all the probability mass in quadrants *II*, *III*, and *IV* within the shaded region. Thus, in this case $F_{HI} = 1 - (1 - \beta)^2$. On the other hand if $\lambda \geq p_{th}$, then $F_{HI} = 1$. Therefore, the optimal choice of λ for a given quantile β that minimizes F_{HI} and maximizes $WPAR$ is $\lambda = \gamma(r_n, \beta)$.

Assuming that the β th quantile for the marginal distribution is the same as that of the nominal Gaussian $\mathcal{N}(\mu(r_n), \sigma^2)$, we have $\gamma(r, \beta) = \mu(r) + \sigma \mathcal{Q}^{-1}(1 - \beta)$. To evaluate the $WPAR$, the nominal fading distribution can be assumed and so:

$$\begin{aligned} P_{FH}(r) &= \mathcal{P}\left(\frac{P_1 + P_2}{2} \leq \lambda | r_{actual} = r\right) \\ &= 1 - \mathcal{Q}\left(\frac{\lambda - \mu(r)}{\frac{\sigma}{\sqrt{2}}}\right) \\ &= 1 - \mathcal{Q}\left(\frac{\mu(r_n) + \sigma \mathcal{Q}^{-1}(1 - \beta) - \mu(r)}{\frac{\sigma}{\sqrt{2}}}\right). \end{aligned} \quad (18)$$

The $WPAR$ can be computed by substituting $P_{FH}(r)$ from (18) into (5).

Figure 12 plots the performance of the averaging detector under the single-quantile model for the fading distribution. The dashed curve (blue) is the performance of the averaging detector when the fading distribution is completely known (in this case the averaging detector is the ML detector). The solid curve (black) is the performance of the averaging detector under minimal knowledge of the fading distribution, i.e., with knowledge of a single quantile. From the figure it is clear that the 2-user averaging detector is highly non-robust to uncertainties in the fading distribution.

The performance of the averaging detector improves if we assume multiple quantile knowledge for shadowing and fading distributions. The mathematical analysis of multiple quantiles is similar to that of the single-quantile model and is omitted here in the interest of space. The performance is shown in Figure 13 and it is clear that the performance of the averaging detector improves as we learn more quantiles about the fading distribution. However, the first few quantiles learned give more performance improvement than the later ones — with performance approaching that of a fully trusted nominal model as the number of trusted quantiles increases.

Although this section's mathematical analysis of the averaging detector covered the two cooperating radios case ($M = 2$), it is fairly straightforward to extend the analysis to $M > 2$. See [43] for the complete details.

Figure 14 shows the performance of the averaging detector under the single-quantile model for $M \geq 2$, with $F_{HI} = 10^{-2}$. The solid curve is the averaging detector with single-quantile knowledge, the dashed-dotted curve is the 'OR-rule' detector (discussed in the next section) and the dashed curve is the averaging detector with complete trust in the nominal distributional (in this case the averaging detector is the ML detector). Note that whereas the performances of the OR




Fig. 12. Averaging detector for two user cooperation: performance under complete knowledge of the fading/shadowing distribution versus performance under the single-quantile uncertainty model for fading/shadowing distribution.




Fig. 13. Performance of the averaging detector under varying quantile knowledge for the shadowing and fading distributions. The quantiles were chosen to minimize F_{HI} for a given $WPAR$ value.

rule and the averaging detector under complete distributional knowledge improve with increasing M , the averaging detector with single-quantile knowledge does worse as the number of cooperating radios M increases! This is because the number of quantiles contributing towards F_{HI} increases exponentially

with the number of users.¹⁹ This shows the non-robustness of blindly using the form of the ML detector under modeling uncertainties.

C. OR-rule detector: hard decision combining

We now explore a more robust detection algorithm that performs well even under minimal models for the fading distribution – the “OR-rule” [63]. This is a hard-decision combining strategy where each radio compares its received power to a threshold. It tentatively declares the band free to use if its received power is below the threshold. Then, each radio sends its tentative 1-bit sensing decision to the central combiner (there are other ways to fuse decision based on the radio topology [85]). The global decision to use the band is made only if all the sensors declare the band to be free.

The safety/performance of the OR rule is easy to compute. Assume that each radio uses a detector threshold $\lambda_{radio,M}$ (detection threshold for a single radio assuming a total of M cooperating radios). Then, the tentative fear of harmful interference for each radio is given by

$$\begin{aligned} F_{HI,radio} &= \mathcal{P}(P_i \leq \lambda_{radio,M} | r_{actual} = r_n) \\ &= 1 - \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r_n)}{\sigma}\right). \end{aligned}$$

This is the same as the fear of harmful interference for the radiometer discussed in Section IV with infinite samples. It is thus clear that a single quantile at the threshold is good enough for a single radio as $N \rightarrow \infty$. The system of cognitive radios causes harmful interference only if every radio individually fails to detect the primary user, and so by the assumption of independence:

$$F_{HI,system} = (F_{HI,radio})^M.$$

In order to meet the target F_{HI} , each radio must choose a $\lambda_{radio,M}$ satisfying

$$\begin{aligned} \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r_n)}{\sigma}\right) &= 1 - [F_{HI}^{target}]^{\frac{1}{M}} \\ \Rightarrow \lambda_{radio,M} &= \sigma \mathcal{Q}^{-1}\left(1 - [F_{HI}^{target}]^{\frac{1}{M}}\right) + \mu(r_n). \end{aligned} \quad (19)$$

As $M \rightarrow \infty$, it is immediately clear that for any given target F_{HI} , the term $[F_{HI}^{target}]^{\frac{1}{M}} \rightarrow 1$ and so the threshold (and thus target quantile) approaches the case of extremely favorable fading.

For such a choice of $\lambda_{radio,M}$, the probability of finding a hole at a radial distance of r is given by

$$\begin{aligned} P_{FH,radio}(r) &= \mathcal{P}(P_i \leq \lambda_{radio,M} | r_{actual} = r) \\ &= 1 - \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r)}{\sigma}\right). \end{aligned}$$

¹⁹To understand why the averaging detector is so vulnerable to uncertainties of this form, remember that the empirical average is very sensitive to outliers. A single very negative number can dominate the entire average. Quantile models can be thought of as histograms. As such, they do not impose any restriction on how negative the rare bad fading can be since the outermost bin of a histogram includes everything from the top of that bin on down to $-\infty$. Consequently, the averaging detector cannot afford even a single user experiencing a fade from that lowermost bin.

The system finds a hole only if all the radios find a hole.

$$\begin{aligned} P_{FH,system}(r) &= (P_{FH,radio})^M \\ &= \left[1 - \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r)}{\sigma}\right)\right]^M \end{aligned} \quad (20)$$

Substituting (20) in (5), we get the WPAR for the OR rule.

It is clear from substituting (19) into (20) that under the single-quantile model of uncertainty and nominal Gaussian fading, the WPAR tends to 1 for the OR rule as the number of cooperating users increases. Figure 12 compares the performance of the OR-rule detector with the averaging detector with complete knowledge, and the averaging detector with single-quantile knowledge for the case of two cooperating radios ($M = 2$) while Figure 14 compares the same for the case of $M > 2$. It is clear that the OR rule is much more robust to uncertainty in the fading distribution than the averaging rule.

Gains by using the OR rule are accomplished by taking the single quantile to correspond to ever more favorable fading realizations. This is problematic since it involves achieving a consensus regarding the rare best fading events — this is as implausible as achieving a consensus regarding the rare worst fading events. In addition, there is a very natural deployment scenario — outdoors on a rooftop — in which the best fading events cannot be too good. This is a little counterintuitive, but remember that multipath fading can result in both destructive and constructive interference. Indoors or in an urban canyon, the best-case fading corresponds to lucky constructive interference. Outdoors, with a dominant line-of-sight path, such constructive interference cannot occur.

Strangely enough, when cooperation is involved, it is this possibility of a clean line-of-sight path that requires the uncertainty model \mathbb{F}_r to impose a bound on how lucky the fading can be. This effectively caps λ_{radio} to the fade that corresponds to a single line-of-sight path. Once the number of cooperating users has reached a point that they can support the desired F_{HI} using that particular quantile, there is no further benefit to increasing the number of users if the OR rule is used. In fact, the performance will drop if cooperating users are blindly added as there is an increased chance of a single user (who happens to be in a rich multipath environment) getting a very lucky constructive fade and thereby deciding that they are within the no-talk radius. The kinked-green curve in Figure 14 illustrates what happens if the uncertain fading model includes the possibility for a line-of-sight path at the 10%-best quantile.

Other weighted-percentage rules for hard-decision combining have also been proposed and these are a little more tolerant of modeling inaccuracies [63] in general. In particular, such rules are required to avoid the performance penalty that arises from the fear of line-of-sight, but there is insufficient space here to discuss them.

D. Performance of cooperation under loss of independence

We have shown that safety/performance can improve significantly if radios cooperatively sense for the primary user as compared to sensing individually. This assumed that the channels from the primary transmitter to the individual secondary radios are independent. However, the primary user might not




Fig. 14. Performance of detectors as a function of the number of cooperating users M for a fixed target $F_{HI} = 10^{-2}$.

trust this assumption since all the cognitive radios may be behind the same obstacle and hence see correlated shadowing. We show that this implies that the detector needs to set its thresholds conservatively, leading to a loss in WPAR performance. The issue of correlated-shadowing is also discussed in [70], where the authors examine the performance of their proposed linear-quadratic detector with correlation uncertainty. The proposed detector is shown to have better probability of detection than a simple counting rule for correlation values greater than 0.4.

As before, let (P_1, P_2, \dots, P_M) be the received powers at the M secondary users. To isolate the effect of dependent shadowing, we assume that the marginal distributions for P_i are completely known, but there is some uncertainty in the correlation across users. For ease of analysis, we assume that (P_1, P_2, \dots, P_M) is a jointly Gaussian random vector with marginals given by $P_i \sim \mathcal{N}(\mu(r), \sigma^2)$, where r is the common radial distance from the primary transmitter. Further, the $M \times M$ covariance matrix C has entries $C(i, j)$ given by

$$C(i, j) = \begin{cases} \rho\sigma^2 & \text{if } i \neq j \\ \sigma^2 & \text{if } i = j \end{cases}$$

where the correlation coefficient ρ is uncertain within known bounds, i.e., $\rho \in [0, \rho_{max}]$, with $0 \leq \rho_{max} \leq 1$.

Under this uncertain correlation model it is easy to show that the averaging detector in (13) is the ML detector no matter what the value of ρ is. Further, it is straightforward to show that to meet a low F_{HI} constraint, the averaging detector must design its λ for the worst case correlation, $\rho = \rho_{max}$. For this choice of ρ we have $\frac{1}{M} \sum_{i=1}^M P_i \sim \mathcal{N}(\mu(r), \frac{1}{M}[1 + (M-1)\rho_{max}]\sigma^2)$. Therefore,

$$F_{HI} = 1 - \mathcal{Q} \left(\frac{\lambda - \mu(r_n)}{\sqrt{\frac{1}{M}[1 + (M-1)\rho_{max}]\sigma^2}} \right).$$

From the above equation we can choose a λ such that the target F_{HI} requirement is met. Given this λ we compute the WPAR performance assuming the nominal model, which corresponds to complete independence, $\rho = 0$, i.e., $\frac{1}{M} \sum_{i=1}^M P_i \sim \mathcal{N}(\mu(r), \frac{1}{M}\sigma^2)$.

Figure 15 shows the performance of the averaging detector designed for different values of ρ_{max} . It is clear that as the amount of uncertainty in the correlation increases, the performance of the averaging detector decreases. Even a small amount of correlation results in a significant drop in performance. As the number of users increases, this particular model of correlation is even more harmful. This can be seen by giving a simple interpretation to this correlation – fading for any user is the sum of a common random fading and a fade local to this user. It is clear that no amount of cooperation can overcome the non-spatially-ergodic common fade. Without a way to combat the fear of such non-spatially-ergodic shadowing uncertainty, there is no way to safely recover the full spectrum hole.




Fig. 15. Performance of the averaging detector with varying amount of correlation uncertainty, ρ_{max} . These plots correspond to the case of ten cooperating users, i.e., $M = 10$.

VI. CALIBRATION AND ASSISTED DETECTION

Section IV analyzed the issues with a single radio trying to achieve a very low F_{HI} . Such a detector must budget for the worst-case multipath and shadowing and hence loses significant area whenever the channel is not badly shadowed. The previous section argued that cooperation among independently shadowed users can help, but offered no hope for the physically important case of users that might experience common shadowing. For example, the uncertainty in the deployment scenario of indoor vs outdoor use can easily manifest as users that are all indoors (and thus shadowed) or outdoors.

If such a radio had information about its shadowing environment it could budget for the actual shadowing and thus improve on its probability of finding a spectrum hole without

giving up any safety. A detection mechanism where side information is used to aid the detector is called calibrated or assisted detection. One example of assisted detection in the cognitive-radio context is interference calibration. If interference from other radios extends beyond a single primary frequency band, then adjacent bands can be used to estimate the interference level and improve the robustness of the detector [16]. Another such example is assisted GPS, where a GPS receiver obtains side information from TV/cell-towers to reduce its uncertainty about its location, time, etc [86], [87]. This section shows how assisted detection improves the $F_{HI}/WPAR$ performance of a single-radio spectrum sensor.

A. Satellite bands as a calibration mechanism

One of the advantages of satellites is that the path loss from a satellite is constant to all places within a large area (for example the San Francisco Bay Area). Hence the signal strength of satellite can be deterministically subtracted to reveal the shadowing + multipath component. So how is satellite shadowing in a separate frequency band related to shadowing from a TV tower in the band of interest? Consider two hypothetical radios: one on the roof of a building and the other in the basement. The radio in the basement will see both the satellite and TV signals at low power levels as compared to the radio on the roof. This suggests that shadowing can be broken up into a directional component that depends on the location of the transmitter and a portion that is direction agnostic. Furthermore, the direction-agnostic shadowing is also wideband – it remains the same across frequencies [81].

B. Satellite-assisted detector

To evaluate the potential gains from using satellite-assisted detection, consider this very simple model.

$$\begin{aligned} P_1 &= p_t - (10 \log_{10}(r^\alpha) + S_1) \\ P_2 &= p_g - (L_g + S_2), \end{aligned} \quad (21)$$

where the new term p_g is the transmit power of the satellite signal, and L_g is the path loss from the satellite transmitter to any radio in this given geographic area. S_1 is the loss due to shadowing and multipath fading encountered by the TV signal, i.e., $S_1 = S + M_1$. Similarly, S_2 is the loss encountered by the satellite signal due to fading and can be written as $S_2 = S + M_2$. We conjecture that the shadowing in the satellite band and the TV band are highly correlated and for simplicity, they are modeled as being identical. M_1, M_2 are independent multipath random variables for the TV and satellite bands. We assume that S_1 and S_2 are normally distributed and are correlated with a correlation coefficient of ρ . Hence,

$$\begin{bmatrix} P_1 \\ S_2 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mu(r) \\ \mu_2 \end{bmatrix}, \begin{bmatrix} \sigma_1^2 & -\rho \sigma_1 \sigma_2 \\ -\rho \sigma_1 \sigma_2 & \sigma_2^2 \end{bmatrix} \right). \quad (22)$$

where $\mu(r) = p_t - (10 \log_{10}(r^\alpha) + \mu_{S_1})$, σ_1^2 is the variance of the primary user's received signal power and σ_2^2 is the variance of the satellite's received signal level. Under this

nominal model, the ML estimate for $\mu(r)$ is given by

$$T(P_1, P_2) = P_1 + \rho \frac{\sigma_1}{\sigma_2} S_2. \quad (23)$$

This test statistic is compared to a threshold to determine if we are inside or outside the no-talk radius:

$$T(P_1, P_2) \stackrel{D=1}{\underset{D=0}{\gtrless}} \lambda. \quad (24)$$

The distribution of the test statistic is given by:

$$P_1 + \rho \frac{\sigma_1}{\sigma_2} S_2 \sim \mathcal{N} \left(\mu_1 + \rho \frac{\sigma_1}{\sigma_2} \mu_2, \sigma_1^2 (1 - \rho^2) \right). \quad (25)$$

Using this distribution (with no additional uncertainty), $F_{HI}(r_n)$ and $P_{FH}(r)$ are:

$$F_{HI} = 1 - \mathcal{Q} \left(\frac{\lambda - (\mu(r_n) + \rho \frac{\sigma_1}{\sigma_2} \mu_2)}{\sigma_1 \sqrt{1 - \rho^2}} \right), \quad (26)$$

$$P_{FH}(r) = 1 - \mathcal{Q} \left(\frac{\lambda - (\mu(r) + \rho \frac{\sigma_1}{\sigma_2} \mu_2)}{\sigma_1 \sqrt{1 - \rho^2}} \right). \quad (27)$$

The performance of this detector compared to a single radio is shown in Figure 16. From the figure, it is evident that the performance improves as the level of correlation between the satellite fading and TV-tower fading increases. This corresponds to when both the signals are wideband and so multipath is relatively less significant. In the extreme of no-multipath ($\rho = 1$), all the area can be recovered by a single satellite-assisted spectrum sensor. There is insufficient space here, but it turns out that if the number of cooperating sensors increases and the non-common shadowing were guaranteed to be independent across sensors, then satellite-assisted techniques can completely overcome the deployment uncertainty that otherwise manifests as the fear of correlated shadowing across users.




Fig. 16. Average case performance of a satellite band assisted receiver versus that of a single radio with no assistance. The variance of satellite band fading was half the variance of the TV fading.

C. Performance of assisted detection with quantile models

The ML detector (weighted-average detector) performs well when the distribution is completely known. Therefore, it is important to consider the effect of uncertainty. For example, suppose that all we knew was that there is only a 1% chance that the satellite shadowing is small (say 5dB) while the TV signal is severely shadowed (say greater than 20dB). How would the weighted energy detector perform with this limited information?

A percentile is chosen for P_1 and another for S_2 (i.e. we choose $(\beta_1, \beta_2, \beta_{12})$ such that $\mathcal{P}(P_1 < \gamma(r, \beta_1)) = \beta_1$, $\mathcal{P}(S_2 < \gamma_g(\beta_2)) = \beta_2$ and $\mathcal{P}(P_1 < \gamma(r, \beta_1), S_2 < \gamma_g(\beta_2)) = \beta_{12}$). Because these are two different bands, there is no reason to assume that the quantiles are the same. Even so, Figure 11 can be used to understand the worst-case performance. The weighted-average detector is $P_1 + \rho \frac{\sigma_1}{\sigma_2} S_2$ which is similar to the diagonal blue line detector in Figure 11, but with a slope that depends on the correlation and relative magnitudes of the fading variances. For the satellite-assisted detector, the probability of region I is $(1 - \beta_1 - \beta_2 + \beta_{12})$. This means that the best value for the threshold λ is given by $\gamma(r_n, \beta_1) + \rho \frac{\sigma_1}{\sigma_2} \gamma_g(\beta_2)$. With this λ the F_{HI} is given by:

$$\begin{aligned} F_{HI}(\beta_1, \beta_2, \beta_{12}, r_n) \\ = \beta_1 + \beta_2 - \beta_{12}. \end{aligned} \quad (28)$$

As before, the WPAR is evaluated using the complete model. The achievable region is the convex hull of all the points generated by changing the values of β_1 and β_2 . This region is shown in Figure 17. The performance is significantly worse than the performance when the channel model is completely trusted. The main reason for the poor performance is that three quadrants $[\mathcal{P}(\{P_1 < \gamma(r, \beta_1)\} \cup \{S_2 < \gamma_g(\beta_2)\})]$ contribute to the F_{HI} for the ML rule.

The counterpart to the OR rule of Section V-C here is the double-threshold detector. This detector declares that the primary user is absent only if $P_1 < \gamma(r, \beta_1)$ and $S_2 < \gamma_g(\beta_2)$ i.e. when the primary signal is low enough and the satellite signal is not significantly faded. The F_{HI} for this detector is β_{12} which is less than the F_{HI} for the weighted-average detector ($\beta_1 + \beta_2 - \beta_{12}$). The performance of this detector is compared to the weighted-average detector in Figure 17. This shows that if the information about the channel model is limited, the double-threshold detector is preferred. In fact, the double-threshold rule with limited knowledge can outperform even a single-radio detector with complete knowledge. This shows that additional information about the shadowing environment is useful even if it is only binary information (i.e. whether we are indoors (deeply shadowed) or outdoors).

VII. CONCLUSIONS

Static frequency planning results in bands being allocated to homogeneous services over large spatial areas and for long times in order to isolate and thus protect the robust operation of heterogeneous wireless systems while preserving their individual freedom. This results in significant underutilization of the spectrum from the perspective of users that could operate on much smaller space-time scales and thereby fit within




Fig. 17. Performance of the double-threshold detector as compared to an ML detector. If the channel model is uncertain, the double-threshold detector is preferred.

the “spectrum holes” left by the static allocations. Dynamic spectrum access can allow the utilization of these spectrum holes.

To do this, strategies for sensing spectrum holes must satisfy two objectives. The first is safety — the primary users must be guaranteed that they will not experience undue interference. The second is efficiency — as much of the spectrum hole as possible must be reclaimed. The core problem is that the incentives of those proposing and implementing the sensing strategy are aligned with the second objective, but not the first. As a result, the primary users have no rational reason to trust the secondary users’ assurances and this results in asymmetric uncertainty.

To reflect this tension and to allow a unified treatment of spectrum sensing, this paper has introduced two distinct metrics. To guarantee safety, the “Fear of Harmful Interference” F_{HI} from the detector must be kept low enough no matter which radio deployment and environmental model turns out to be true. It is only the primary user’s uncertainty set that matters here. Consensus between the regulators, primary users, and secondary users has to be achieved regarding this uncertainty and so it is likely to remain large. Every sensing strategy will have its own critical uncertainties that must be bounded.

In this paper, quantile models have been proposed for uncertain probability distributions (e.g. for shadowing) while secondary radio positions have been considered unconstrained. In multiuser settings, the degree of shadowing correlation turns out to be a very significant uncertainty. It has been suggested that it might be easier to achieve a firm consensus regarding the correlation of shadowing across different frequencies for a single radio than it is to achieve a consensus regarding the shadowing correlation across users. This remains to be explored more fully, but the F_{HI} metric seems to be the right way to capture the otherwise vague notion of safety while still

allowing significant innovation at the detector level.

For performance purposes, the secondary user has no reason to lie and is free to analyze its own performance using any desired probability model and utility function. The core issue here is one of simplicity and generalizability. To enable high-level comparisons between detection strategies, it is important to be able to decouple the interaction among different primary users while capturing the key effects. Restricting attention to spectrum holes that are very long lived in time, we have argued that the most significant terms are:

- As we get further away from any primary transmitter, there is an increasing chance that we will be within the service footprint of another primary transmitter.
- Area closer to the primary user’s footprint is more valuable (in a business or utility sense) than area far away because primary users are likely to have positioned their transmitters so as to serve a maximal number of humans.

Therefore, the proposed “Weighted Probability of Area Recovered” (WPAR) metric uses a discounting-function to weigh the probability of recovering area at a given distance away from a single primary transmitter. While exponential discounting has been used here for convenience, it remains an interesting open question to determine what the right discounting functions are for different application scenarios.

By using these metrics, this paper has shown that the popularly used metrics of sensitivity, P_{FA} , and P_{MD} (such as the -116dBm rule used by the IEEE 802.22 process) are overly constraining. Even an ideal detector (one with $P_{FA} = P_{MD} = 0$ for a desired sensitivity) has poor WPAR performance when facing uncertain fading. Too much valuable area must be sacrificed to achieve the desired robustness — effectively turning this into a static guard band by another name. However, the F_{HI} and WPAR metrics allow the principled consideration of alternative strategies such as multiuser cooperation and show exactly which uncertainties must be resolved (and to what resolution) in order to be able to guarantee both safety and high performance for a detector of a given complexity. Therefore, we suggest that specifications for detection strategies be expressed at the F_{HI} and WPAR level rather than in terms of a desired sensitivity and ROC.

This paper represents the beginning of a story rather than the end of one. Much remains to be done. In particular:

- Cooperative sensing strategies that utilize assisted detection need to be analyzed. The performance of such strategies under our new F_{HI} and $WPAR$ metrics needs to be evaluated [88].
- The tradeoffs between the time-overhead (sensing time + cooperative message exchange) and the space-overhead (WPAR effects + sensing-MAC effects [89]) need to be understood. It is here that different signal-processing strategies are likely to distinguish themselves. Such a space-time hole is illustrated in Figure 18 where the combination of temporal and spatial margins/overheads is illustrated.
- Under the traditional metric of sensitivity, an SNR wall for a sensing algorithm sets a bound on how sensitive a detector can be given the uncertain model for the

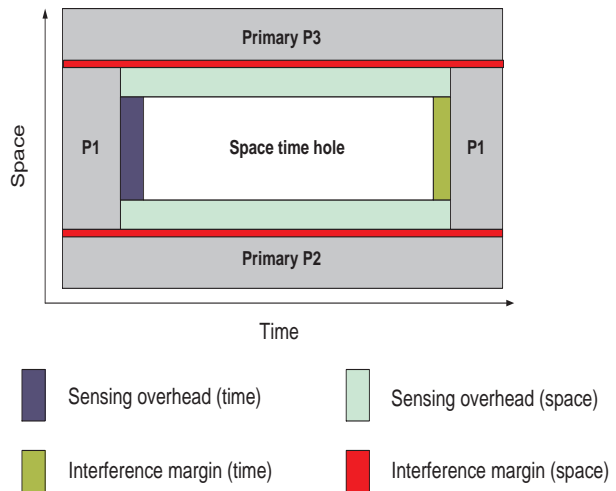


Fig. 18. A space-time hole. Primary users P_1 , P_2 and P_3 occupy different space-time regions but the same frequency band. The secondary user can recover the hole when P_1 ceases transmission albeit with some time lost due to temporal sensing overhead. When P_1 reappears, the secondary user can still transmit for a finite duration (temporal interference margin). Corresponding spatial interference margins and sensing overheads are also shown in the figure.

noise process [16]. The role of SNR walls must also be understood in the context of WPAR and F_{HI} since sensitivity is now implicit rather than explicit.

- The simple quantile models that have been proposed here are intuitively clear and easy to use but clearly do not represent the form of uncertainty representation that is both unambiguously verifiable and realistic. The example of the subtle role of constructive interference in Section V-C made that clear. Since consensus is required between primary and secondary users, one would prefer an uncertainty model that came with a experimental certificate of correctness.
- The F_{HI} metric currently captures only one dimension of fear — that of optimistic assumptions regarding the environment. In practice, there is also the fear of dishonest implementations.²⁰ The regulators, primary users, and secondary users should only need to achieve consensus regarding some key features of the wireless system implementation rather than for every aspect.²¹ The safety of the rest of the implementation should rely on self-regulation (or peer regulation) through the design of an appropriately lightweight enforcement mechanism.

ACKNOWLEDGMENTS

We thank the United States National Science Foundation (ANI-326503, CNS-403427, CCF-729122), C2S2 (Center for Circuit System Solutions), and Sumitomo Electric for their generous funding support. Students Kristen Ann Woyach,

²⁰Even after we agree to the rules of a game, we still need a way to make sure that players adhere to those rules [90].

²¹Some might wonder “why can’t we all just get along?” It is important to remember that there are business reasons why primary users might wish to hinder secondary use so as to forestall competition [91]. Thus it is preferable to have only a few such regulatory hurdles.

Pulkit Grover, and Hari Palaiyanur are thanked for their careful comments on drafts of this paper. The anonymous reviewers and editors were also very helpful in sharpening the paper's presentation and provided some important references.

REFERENCES

- [1] "Establishment of an interference temperature metric to quantify and manage interference and to expand available unlicensed operation in certain fixed, mobile and satellite frequency bands," Federal Communications Commission, Tech. Rep. 03-237, Nov. 2003. [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-03-289A1.pdf
- [2] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE J. Select. Areas Commun.*, vol. 23, pp. 201–220, Feb. 2005.
- [3] "Spectrum policy task force report," Federal Communications Commission, Tech. Rep. 02-135, Nov 2002. [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-228542A1.pdf
- [4] M. A. McHenry and K. Steadman, "Spectrum occupancy measurements, location 1 of 6: Riverbend park, Great Falls, Virginia," Tech. Rep., 2005. [Online]. Available: http://www.sharedspectrum.com/inc/content/measurements/nsf/1_NSF.Riverbend.Park.Report.pdf
- [5] 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
- [6] 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, pp. 1662–1671.
- [7] N. Hoven and A. Sahai, "Power scaling for cognitive radio," in *Proc. of the WirelessCom 05 Symposium on Emerging Networks, Technologies and Standards*, Maui, HI, June 2005, pp. 13–16.
- [8] N. Hoven, "On the feasibility of cognitive radio," Master's thesis, University of California, Berkeley, 2005. [Online]. Available: http://www.eecs.berkeley.edu/~sahai/Theses/Niels_MSThesis.pdf
- [9] M. Maddah-Ali, A. Motahari, and A. Khandani, "Communication over MIMO X-channels: interference alignment, decomposition, and performance analysis," *IEEE Trans. Inform. Theory*, vol. 54, pp. 3457–3470, Aug. 2008.
- [10] V. R. Cadambe and S. A. Jafar, "Interference alignment and the degrees of freedom for the k user interference channel," *IEEE Trans. Inform. Theory*, vol. 54, pp. 3425–3441, Aug. 2008.
- [11] S. Haykin, "Cognitive dynamic systems," *Proc. IEEE*, vol. 94, pp. 1910–1911, Nov. 2006.
- [12] I. J. Mitola, "Software radios: Survey, critical evaluation and future directions," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 8, pp. 25–36, Apr. 1993.
- [13] A. Ghasemi and E. S. Sousa, "Spectrum sensing in cognitive radio networks: requirements, challenges and design trade-offs [cognitive radio communications]," *IEEE Commun. Mag.*, no. 4, pp. 32–39, Apr. 2008.
- [14] S. Shellhammer, S. Shankar, R. Tandra, and J. Tomcik, "Performance of Power Detector Sensors of DTV Signals in IEEE 802.22 WRANs," in *Proc. of the First ACM International Workshop on Technology and Policy for Accessing Spectrum (TAPAS)*, Aug. 2006. [Online]. Available: <http://doi.acm.org/10.1145/1234388.1234392>
- [15] Y. Li, "Signal processing issues in Cognitive Radio," *Proc. IEEE*, 2008.
- [16] R. Tandra and A. Sahai, "SNR walls for signal detection," *IEEE Journal on Selected Topics in Signal Processing*, vol. 2, pp. 4 – 17, Feb. 2008.
- [17] M. Ghosh, "Text on FFT-based pilot sensing - For Informative Annex on Sensing Techniques," IEEE 802.22 Meeting Documents, July 2007, <https://mentor.ieee.org/802.22/file/07/22-07-0298-01-0000-text-on-fft-based-pilot-sensing.doc>
- [18] D. Cabric, A. Tkachenko, and R. W. Brodersen, "Spectrum sensing measurements of pilot, energy and collaborative detection," in *Military Communications Conference*, Washington, DC, Oct. 2006, pp. 1–7.
- [19] S. Shellhammer and R. Tandra, "An evaluation of DTV pilot power detection," in *IEEE 802.22-06/0188r0*, Sept. 2006. [Online]. Available: http://ieee802.org/22/Meeting_documents/2006_Sept/22-06-0188-00-0000-An-Evaluation-of-DTV-Pilot-Power-Detection.ppt
- [20] R. Tandra and A. Sahai, "Noise calibration, delay coherence and SNR walls for signal detection," in *Proc. of the 3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.
- [21] W. A. Gardner, "Signal interception: A unifying theoretical framework for feature detection," *IEEE Trans. Commun.*, vol. 36, pp. 897–906, Aug. 1988.
- [22] —, "Signal interception: Performance advantages of cyclic-feature detectors," *IEEE Trans. Commun.*, vol. 40, pp. 149–159, Jan. 1992.
- [23] H. S. Chen and W. Gao, "Text on Cyclostationary Feature Detector - For Informative Annex on Sensing Techniques," IEEE 802.22 Meeting Documents, July 2007, <https://mentor.ieee.org/802.22/file/07/22-07-0283-00-0000-text-on-cyclostationary-feature-detector-thomson.doc>
- [24] S. Haykin, J. Reed, and D. Thomson, "Spectrum sensing in Cognitive Radio," *Proc. of the IEEE*, 2008.
- [25] R. Tandra and A. Sahai, "SNR walls for feature detectors," in *Proc. of 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 2007, pp. 559–570.
- [26] G. Trukenich, V. Gaddam, and M. Ghosh, "Text on Dual FPLL pilot sensing - For Informative Annex on Sensing Techniques," IEEE 802.22 Meeting Documents, July 2007, <https://mentor.ieee.org/802.22/file/07/22-07-0296-00-0000-text-on-dual-fpll-pilot-sensing.doc>
- [27] Y. Zeng and Y.-C. Liang, "Text on eigenvalue based sensing - For Informative Annex on Sensing Techniques," IEEE 802.22 Meeting Documents, July 2007, <https://mentor.ieee.org/802.22/file/07/22-07-0297-03-0000-text-on-eigenvalue-based-sensing.doc>
- [28] —, "Maximum-Minimum Eigenvalue Detection for Cognitive Radio," in *Proc. of the IEEE 18th International Symposium on Personal, Indoor and Mobile Radio Communications, (PIMRC)*, Sept. 2007, pp. 1–15.
- [29] V. V. Veeravalli, "Decentralized quickest change detection," *IEEE Trans. Inform. Theory*, vol. 47, no. 4, pp. 1657–1665, May 2001.
- [30] A. Parsa, A. A. Gohari, and A. Sahai, "Exploiting interference diversity for event-based spectrum sensing," in *Proc. of the 3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, 2008.
- [31] S. M. Kay, *Fundamentals of statistical signal processing: Detection theory*. Upper Saddle River, New Jersey: Prentice Hall PTR, 1998, vol. 2.
- [32] Z. Quan, S. Cui, A. H. Sayed, and H. V. Poor, "Wideband Spectrum Sensing in Cognitive Radio Networks," in *Proc. of the IEEE International Conference on Communications (ICC)*, May 2008, pp. 19–23.
- [33] Y.-C. Liang, Y. Zeng, E. C. Peh, and A. T. Hoang, "Sensing-Throughput Tradeoff for Cognitive Radio Networks," *IEEE Trans. Wireless Commun.*, no. 4, pp. 1326–1337, Apr. 2008.
- [34] S. Geirhofer, L. Tong, and B. M. Sadler, "Dynamic spectrum access in the time domain: Modeling and exploiting white space," *IEEE Commun. Mag.*, vol. 45, pp. 66–72, May 2007.
- [35] —, "A measurement-based model for dynamic spectrum access in WLAN channels," in *Military Communications Conference*, Washington, DC, Oct. 2006, pp. 1–7.
- [36] Q. Zhao, S. Geirhofer, L. Tong, and B. M. Sadler, "Optimal dynamic spectrum access via periodic channel sensing," in *IEEE Wireless Communications and Networking Conference*, Hong Kong, Mar. 2007, pp. 33–37.
- [37] V. Paxson and S. Floyd, "Wide area traffic: the failure of Poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, pp. 226–244, 1995.
- [38] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol I*, 3rd ed. Belmont, Massachusetts: Athena Scientific, 2007.
- [39] —, *Dynamic Programming and Optimal Control, Vol II*, 3rd ed. Belmont, Massachusetts: Athena Scientific, 2007.
- [40] J. Huang, R. A. Berry, and M. L. Honig, "Distributed interference compensation for wireless networks," *IEEE J. Select. Areas Commun.*, vol. 24, no. 5, pp. 1074–1084, May 2006.
- [41] K. Eswaran, M. Gastpar, and K. Ramchandran, "Bits through ARQs: Spectrum sharing with a primary packet system," in *IEEE International Symposium on Information Theory*, Nice, France, June 2007, pp. 2171–2175.
- [42] C. R. Stevenson, C. Cordeiro, E. Sofer, and G. Chouinard, "Functional requirements for IEEE 802.22 WRAN standard," Tech. Rep., Sept. 2005. [Online]. Available: <https://mentor.ieee.org/802.22/file/05/22-05-0007-48-0000-draft-wran-rqmts-doc.doc>
- [43] R. Tandra, S. M. Mishra, and A. Sahai, "What is a spectrum hole and what does it take to recognize one: extended version," University of California, Berkeley, Berkeley, CA, Tech. Rep., Aug. 2008.
- [44] S. M. Mishra, R. Tandra, and A. Sahai, "Coexistence with primary users of different scales," in *Proc. of 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 2007, pp. 158–167.

- [45] B. Wild and K. Ramchandran, "Detecting primary receivers for cognitive radio applications," in *Proc. of 1st IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, USA, Nov. 2005, pp. 124–130.
- [46] N. Devroye, P. Mitran, and V. Tarokh, "Achievable rates in cognitive radio channels," *IEEE Trans. Inform. Theory*, vol. 52, pp. 1813–1827, May 2006.
- [47] A. Jovicic and P. Viswanath, "Cognitive radio: An information-theoretic perspective," in *IEEE International Symposium on Information Theory*, Seattle, USA, July 2006, pp. 2413–2417.
- [48] P. Grover and A. Sahai, "On the need for knowledge of the phase in exploiting known primary transmissions," in *Proc. of 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 2007, pp. 462–471.
- [49] R. Tandra and A. Sahai, "Is interference like noise when you know its codebook?" in *IEEE International Symposium on Information Theory*, Seattle, USA, July 2006, pp. 2220–2224.
- [50] "Method for point-to-area predictions for terrestrial services in the frequency range 30 MHz to 3000 MHz," ITU-R, Tech. Rep. P.1546-3, Jan. 2007.
- [51] S. J. Shellhammer, S. N. Shankar, R. Tandra, and J. Tomcik, "Performance of power detector sensors of DTV signals in IEEE 802.22 WRANs," in *First International Workshop on Technology and Policy for Accessing Spectrum*, June 2006.
- [52] N. A. Robert Price, "Detection theory," *IEEE Trans. Inform. Theory*, vol. 7, pp. 135–139, July 1961.
- [53] G. Chouinard, "DTV signal stochastic behavior at the edge of the protected contour and resulting probability of detection from various sensing schemes," IEEE 802.22 Meeting Documents, March 2007. [Online]. Available: http://www.ieee802.org/22/Meeting_documents/2007_Mar/22-07-0112-00-0000_Sensing_Stochastic_Behavior.ppt
- [54] P. Samuelson, "A note on measurement of utility," *The Review of Economic Studies*, vol. 4, no. 2, pp. 155–161, Feb 1937.
- [55] D. A. Ryland, "Sutton's or Dock's law?" *New England Journal of Medicine*, vol. 302, no. 17, p. 972, Apr. 1980.
- [56] H. Niedercorn, "A Negative Exponential Model of Urban land use densities and its implications for Metropolitan Development," *Journal of Regional Science*, vol. 11, pp. 317–326, May 1971.
- [57] J. Hirshleifer and J. G. Riley, *The Analytics of Uncertainty and Information*, ser. Cambridge Surveys of Economic Literature Series, 1992.
- [58] F. Database, "Summary of channel 30 transmitters." [Online]. Available: <http://www.fcc.gov/fcc-bin/tvq?state=&call=&arn=&city=&chan=30&cha2=30&serv=&type=0&facid=&list=2&dist=&dlat2=&mlat2=&slat2=&dlon2=&mlon2=&slon2=&size=9>
- [59] J. Notor, "Personal communication," 2004.
- [60] A. Sahai and D. Cabric, "Cyclostationary feature detection," Tutorial presented at DYSPAN, Baltimore, MD, October 2005. [Online]. Available: http://www.eecs.berkeley.edu/~sahai/Presentations/DySPAN05_part2.ppt
- [61] A. Sendonaris, E. Erkip, and B. Aazhang, "User Cooperation Diversity—Part I: System Description," *IEEE Trans. Commun.*, vol. 51, no. 11, pp. 1927–1938, November 2003.
- [62] P. K. Varshney, *Distributed Detection and Data Fusion*. Springer-Verlag, 1997.
- [63] S. M. Mishra, A. Sahai, and R. W. Brodersen, "Cooperative sensing among cognitive radios," in *IEEE International Conference on Communications*, vol. 4, June 2006, pp. 1658–1663.
- [64] G. Ganesan and Y. Li, "Cooperative spectrum sensing in cognitive radio, part I: Two user networks," *IEEE Trans. Wireless Commun.*, pp. 2204–2213, June 2007.
- [65] —, "Cooperative spectrum sensing in cognitive radio, part II: Multiuser networks," *IEEE Trans. Wireless Commun.*, pp. 2214–2222, June 2007.
- [66] T. Weiss, J. Hillenbrand, and F. Jondral, "A diversity approach for the detection of idle spectral resources in spectrum pooling systems," in *Proc. of the 48th Int. Scientific Colloquium, Ilmenau, Germany*, 2003.
- [67] E. Vistotsky, S. Kuffner, and R. Peterson, "On Collaborative Detection of TV Transmissions in Support of Dynamic Spectrum Sharing," in *Proc. of 1st IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 338–345.
- [68] A. Ghasemi and E. S. Sousa, "Collaborative Spectrum Sensing for Opportunistic Access in Fading Environments," in *Proc. of 1st IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Baltimore, MD, Nov. 2005, pp. 131–136.
- [69] J. Ma and Y. G. Li, "Soft combination and detection for cooperative spectrum sensing in cognitive radio networks," in *Proc. of the IEEE Global Telecommunications Conference (GLOBECOM)*, Nov. 2007, pp. 3139–3143.
- [70] J. Unnikrishnan and V. Veeravalli, "Cooperative sensing for primary detection in cognitive radios," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, pp. 18–27, Feb. 2008.
- [71] R. Chen, J.-M. Park, and K. Bian, "Robust Distributed Spectrum Sensing in Cognitive Radio Networks," Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Technical Report TR-ECE-06-07, July 2006.
- [72] R. Chen, J.-M. Park, and J. Reed, "Defense against Primary User Emulation Attacks in Cognitive Radio Networks," *IEEE J. Select. Areas Commun.*, vol. 26, pp. 25–37, Jan. 2008.
- [73] P. Kaligineedi, M. Khabbazi, and V. Bhargava, "Secure Cooperative Sensing Techniques for Cognitive Radio Systems," in *Proc. of the IEEE International Conference on Communications (ICC)*, May 2008, pp. 3406–3410.
- [74] C. Song and Q. Zhang, "Sliding-Window Algorithm for Asynchronous Cooperative Sensing in Wireless Cognitive Networks," in *Proc. of the IEEE International Conference on Communications (ICC)*, May 2008, pp. 3432–3436.
- [75] N. Ahmed, D. Hadaller, and S. Keshav, "GUESS: gossiping updates for efficient spectrum sensing," in *Proceedings of the 1st international workshop on Decentralized resource sharing in mobile computing and networking*, 2006, pp. 12–17.
- [76] G. Ganesan, Y. Li, B. Bing, and S. Li, "Spatiotemporal Sensing in Cognitive Radio Networks," *IEEE J. Select. Areas Commun.*, vol. 26, pp. 5–12, Jan. 2008.
- [77] X. Huang, N. Han, G. Zheng, S. Sohn, and J. Kim, "Weighted-Collaborative Spectrum Sensing in Cognitive Radio," in *Second International Conference on Communications and Networking in China (CHINACOM)*, Aug. 2007, pp. 110–114.
- [78] W. Zhang, R. Mallik, and K. B. Letaief, "Cooperative Spectrum Sensing Optimization in Cognitive Radio Networks," in *Proc. of the IEEE International Conference on Communications (ICC)*, May 2008, pp. 3411–3415.
- [79] F. Visser, G. Janssen, and P. Pawelczak, "Multinode Spectrum Sensing Based on Energy Detection for Dynamic Spectrum Access," in *Proc. of the IEEE Vehicular Technology Conference (VTC)*, May 2008, pp. 1394–1398.
- [80] S. Mishra and R. W. Brodersen, "Cognitive Technology for improving Ultra-Wideband (UWB) Coexistence," in *Proc. of the IEEE International Conference on Ultra-Wideband (ICUWB'07)*, Sept. 2007, pp. 253–258.
- [81] S. Mishra, R. Tandra, and A. Sahai, "The case for Multiband Sensing," in *Proc. of the Forty-fifth Annual Allerton Conference on Communication, Control, and Computing*, Sept. 2007.
- [82] G. Atia, E. Ermis, and V. Saligrama, "Robust Energy Efficient Cooperative Spectrum Sensing in Cognitive Radios," in *Proc. of the IEEE/SP 14th Workshop on Statistical Signal Processing, (SSP'07)*, Aug. 2007, pp. 502–506.
- [83] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*, 1st ed. Cambridge, United Kingdom: Cambridge University Press, 2005.
- [84] M. Gudmundson, "Correlation model for Shadow fading in Mobile Radio Systems," *Electronic Letters*, vol. 27, no. 23, pp. 2145–2146, 1991.
- [85] J. N. Tsitsiklis, *Advances in Signal Processing*, Vol. 2. JAI Press, 1993, ch. Decentralized Detection, pp. 297–344.
- [86] Qualcomm Inc., "The Qualcomm/SnapTrack wireless-assisted GPS hybrid positioning system and results from initial commercial deployments," Tech. Rep., Oct. 2002.
- [87] N. Agarwal *et al.*, "Algorithms for GPS operation indoors and downtown," *GPS Solutions*, no. 6, pp. 149–160, 2002.
- [88] S. M. Mishra, R. Tandra, and A. Sahai, "Calibration/assisted detection for cognitive radios," *IEEE Trans. Signal Processing*, 2008, in progress.
- [89] C. Cordeiro, M. Ghosh, D. Cavalcanti, and K. Challapali, "Spectrum sensing for dynamic spectrum access of TV bands," in *Second International Conference on Cognitive Radio Oriented Wireless Networks and Communications*, Orlando, Florida, Aug. 2007, pp. 225–233.
- [90] K. A. Woyach, A. Sahai, G. Atia, and V. Saligrama, "Crime and punishment for cognitive radios," in *Forty-sixth Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sept. 2008.
- [91] N. Isaacs, "Barrier activities and the courts: A study in anti-competitive law," *Law and Contemporary Problems*, vol. 8, no. 2, pp. 382–390, 1941.