Exploiting Interference Diversity for Event-Based Spectrum Sensing

Arash Parsa, Amin Aminzadeh Gohari and Anant Sahai
Wireless Foundations and Berkeley Wireless Research Center
Department of Electrical Engineering and Computer Sciences
University of California at Berkeley
Emails: {parsa, aminzade, sahai}@eecs.berkeley.edu

***Abstract*—Spectrum sensing is a core problem in cognitive radio. Detecting the presence/absence of very weak primary users with a single antenna can be very difficult. Environmental uncertainties result in SNR walls that detectors cannot overcome. Multiple-antenna approaches show some potential gains, but we show here that single-user multiple-antenna detection still must suffer from an SNR wall. The reason is that the real-world uncertainty in noise is dominated by the potential presence of an unknown number of low-power and time-varying interference sources in the external environment.**

Traditional collaborative spectrum sensing attempts to use the shadowing/multipath diversity across different users to boost the reliability of detection. We show that there is another kind of diversity that is also available: interference diversity. This captures the fact that low-powered interference sources are local to individual users whereas the primary user has a global footprint. To exploit this diversity, we must shift our perspective from existence-based detection (whether the primary is present or not) to event-based detection (whether the primary has turned off or on).

I. Introduction

Spectrum sensing at very low SNR regimes has gained crucial importance in the context of opportunistic spectrum sharing. For example, the required sensitivity in the IEEE 802.22 standard is -116 dBm (-22 dB of signal to noise ratio) [1]. In addition to high sensitivity, the secondary users need to be able to operate in different settings and environments. Uncertainty in the noise environment imposes fundamental limits on the sensitivity of secondary users. Below a certain signal-to-noise ratio (SNR), detectors will not be able to reliably sense the existence of the signal no matter how long they collect data [2]. This level is called the SNR wall of the corresponding detector and partially quantifies the robustness of the sensing system.

For the sensing problem, different algorithms such as energy detection, matched filtering and cyclostationary feature detection have been proposed [3], [4], [5], [6]. Table I shows a categorization of them by the number of users/antennas. Implicit in the traditional discussion is the nature of the spectral hole that the secondary radios are trying to find. As pointed out in [11], spectrum holes are regions in space-time-frequency within which non-interfering use is possible for a secondary user. Many of the currently used algorithms choose to test the existence/absence of the primary signal. We refer to these detectors as "existence-based" detectors. These hypotheses implicitly assume that the focus is on the spatial-dimension of the spectrum hole (are we close enough to interfere with a primary user).

TABLE I
List of a few noncoherent detection algorithms

Example Detectors	Characteristics
Radiometer	single-antenna; non-cooperative
Max-Min Eigenvalue [7]	multi-antenna; non-cooperative
[8], [9], [10] **The proposed detector**	multi-antenna cooperative

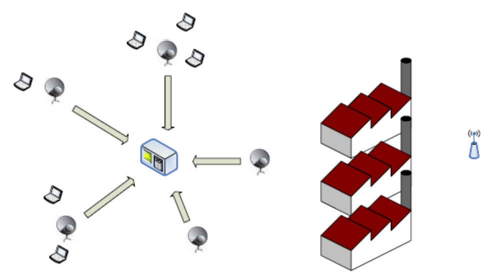

Fig. 1. Interference Diversity: unintentional interferers (e.g. laptops in the figure) are sensed locally while the primaries are sensed globally.

However, for any given secondary-user, it is a time-centered perspective that frames the question in steady state: is it safe to keep using this channel or should we start using that channel? The user's own physical motion through space-time connects the global space-time perspective to a local timeline. The purely spatial perspective on sensing spectrum holes thus corresponds to either the start-up transient (when time is beginning for a user) or to the case of very rapid motion (approximated by a user that can randomly teleport between sensing times). Given that the primary footprints for digital-television stations are on the scale of kilometers and many secondaries have motion only in the meters/second range with reasonable sensing times of less than a second, it seems natural to take a closer look at what the right hypotheses really are in steady-state.

Assuming that secondary users are relatively stationary with respect to primary footprints, we see that the time-centered perspective offers a new set of hypotheses: the entrance/exit/no-change of the primary signal. Detectors that attempt to segment time based on these are called "event-based" detectors. It turns out that this choice of perspective greatly influences the structure and performance of the corresponding hypothesis-testing algorithms.

This has been viewed in the context of "change detection" algorithms. Veeravalli [12] and Mei [13] study the application of quickest-detection theory to decentralized decision systems. Quickest detection [14], [15] refers to real-time detection of abrupt changes in the distribution of observed signals as quickly as possible. Furthermore, Li *et. al.* [16] apply the theory of quickest detection to the problem of sensing primary user signals. However, the previous approaches do not capture the problem of noise uncertainty and robustness. In particular, the real-world presence of time-varying interference can cause distributional changes of the received signal even if there is no primary activity.

Table II compares existence-based and event-based detection in the context of the energy detector.[1] The interesting fact is that it is interference that is the fundamental cause of sensing limitations. Interference is some combination of "unintentional emitters" (e.g. a laptop leaking power outside its permitted bandwidth during a communication, or the electronics of a printer radiating electromagnetic waves during its activity) and "intentional emitters" (other secondary networks using the bandwidth when the primary is absent). A sensor may have some knowledge about the activity patterns of the intentional emitters. If these are other secondary users, this uncertainty can be reduced by using a "sensing-MAC protocol" that commands nearby secondaries to be quiet while one is sensing [4], [17]. However, there is no way of even interacting with unintentional emitters. Furthermore, it is implausible to believe that unintentional emitters will comply with any probabilistic model.

The key insight in this paper is that we know from physical considerations that such unintentional interference is low powered and hence local in the sense that the arrival or departure of one such interferer can only cause a few nearby sensors to trigger false alarms. In contrast, when a primary signal enters or exits, it does so with a global footprint causing many sensors to fire alarms simultaneously. This suggests that a cooperative multi-sensor algorithm can exploit what we call "interference diversity" to distinguish between the two events by gathering information from sensors placed at different physical locations (see Figure 1). In fact, as long as the above local/global property holds, the proposed algorithm succeeds.[2]

This paper is organized as follows. To make this paper self-contained, Section II reviews a couple of existence-based sensing algorithms. The role of external interferers in imposing fundamental limits on the performance of this class of algorithms is highlighted — in particular an SNR wall is discussed for the max-min-eigenvalue detector with the proof deferred to Appendix A. Section III introduces the event-based cooperative sensing algorithm and it is analyzed at an intuitive level, deferring the rigor-

[1]The same ideas can be applied to other detection algorithms such as matched filtering and cyclostationary feature detection at the expense of notational complexity.

[2]Remember, all this holds only if we accept that the cooperating users are within the same spatial spectrum hole and stay there. There are further considerations when secondary users can move long distances. For instance, assume that there are two TV towers, one in Sacramento and another in San Francisco. If we have a network of secondary users close to the Sacramento area, it will be able to correctly detect the entrance/exit of the primary transmitter's signal there. Now, suppose that a few secondary users move to the San Francisco area. Even if they correctly detect the San Francisco tower's transmission, since they are few in number, they will not be able to change the decision made by the whole network, causing the system to fail. Hence, interference diversity is no substitute for the basic requirement of all cooperative algorithms: that all the nodes cooperating to sense are indeed trying to sense the same spectral hole! However, it is interesting to see that unlike cooperating for existence-based sensing, there is no requirement that the collective footprint of the cooperating secondary users be much smaller than the primary user's service area footprint. In fact, having some cooperating users very close to the potentially-on primary user while others are very far away is even desirable!

TABLE II
COMPARISON OF ENERGY-BASED DETECTION IN THE EXISTENCE AND EVENT FRAMEWORKS

Detector:	**Traditional Existence-based**	**Event-based**
Hypotheses:	High-power vs low-power	Increased vs decreased vs unchanged power
Detector statistic:	Received power-level	Change in received power-level
False alarms:	Noise higher than expected	Noise increases suddenly
Missed detections:	Noise lower than expected	Noise drops suddenly
Dom. uncertainty:	The number of active interferers	The activity pattern of interferers

ous formulation and analysis to Appendix B. Section IV includes some simulation results.

II. EXISTENCE-BASED DETECTORS

In this section, we briefly review detection that is based on the following two hypotheses:

$\mathbf{H}_0$: only noise $x[n] = w[n]$
$\mathbf{H}_1$: signal + noise $x[n] = w[n] + r[n]$

where $x[n]$, $w[n]$ and $r[n]$ are the received signal, the input noise and the primary signal at time n. We refer to these detectors as "existence based" or $H_0 - H_1$ detectors. The sensitivity is how weak of a $r[n]$ signal can still be reliably detected or ruled-out.

In general, by increasing the length of the observation window, a sensing algorithm will have a better sensitivity. However, when there are uncertainties in the system, SNR walls exist that impose a fundamental limitation to the sensitivity even if the observation window becomes infinitely long [2]. As the previous section argued, unintended emitters are an irreducible system uncertainty. The goal of the following subsections is to highlight the impact of unknown interferers on detection algorithms at low SNRs. For illustrative purposes, the energy detector (radiometer) is used in the single-antenna context while the maximum-minimum eigenvalue detector is used in the multiple-antenna context.

A. Single-Antenna Energy Detection: Radiometer

An energy detector (radiometer) measures the average received power in a time window, and makes a decision by comparing it to a fixed threshold. If the noise variance were perfectly known, then an increased observation window N would lead the detector to be able to decide between the hypotheses correctly with arbitrarily high probability.

The noise uncertainty for the radiometer is modeled by assuming that the noise variance is known only up to an interval. The length of this uncertainty interval determines the "SNR wall" for the radiometer (see Figure 2) [2]. It is easy to see that the radiometer's SNR wall can

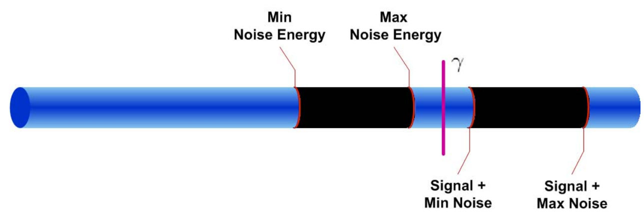

Fig. 2. The noise variance has a range. When the primary signal is present, this interval shifts. If the signal is weak and these two intervals collide, there is an ambiguity that cannot be resolved.

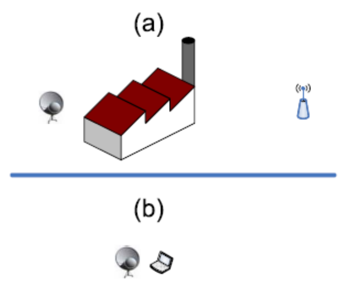

Fig. 3. Single-antenna radiometer fails to distinguish a weak primary signal from a local interferer.

never be less than the **maximum** number of significant interferers that might co-exist in an environment times the average power of a typical significant interferer (see Figure 3).

B. Max-Min Eigenvalue Detection

Liang and Zeng in [7] introduced the multi-antenna sensing algorithm known as the max-min eigenvalue detector. Their core assumption is that noise is white and uniform (in terms of power) over different antennas. The authors argue that if the primary signal is absent, the correlation between samples at different antennas and/or different times will be zero. If the primary signal is present, some of these correlations will not be zero. Therefore the detector can distinguish between these two scenarios by measuring the deviation of a covariance matrix from the identity matrix in terms of the ratio of the maximum and minimum eigenvalues. If the assumption were true, the max-min eigenvalue detector would overcome any noise uncertainty.

Unfortunately, local unintended emitters are an irreducible source of uncertainty and they do not satisfy this assumption. Having a physically localized source, their interference will generally not be white and uniform (in terms of power) over different antennas. From the perspective of the max-min eigenvalue detector, there is no difference between a very weak primary and a local emitter. Appendix A contains a proof for a theorem that makes all this precise.

III. Event-Based Detection

As discussed in the previous section, sensing algorithms based on existence/absence ($H_0 - H_1$) hypotheses are vulnerable to external interferers (which in effect masquerade as weak primaries). In what follows, we make the assumption that the secondary radio is fixed in position relative to the primary transmitter. We further make the assumption that we are operating in steady-state[3] and are thus interested in the time-dimension of the spectrum holes [11]. For simplicity of exposition, we focus on an event-oriented variant of the radiometer.

Intuitively, the transition between the on/off states of the primary signal creates rapid changes in the received power at the spectrum sensors. We refer to these changes as "energy edges"; if there is a upward surge in the power level, we call it a "positive edge" while a downward surge is called a "negative edge". For every time instance, we consider three event hypotheses:

$$\begin{cases} \mathbf{H}_-[n]: & \text{there is a negative edge at time } n; \\ \mathbf{H}_+[n]: & \text{there is a positive edge at time } n; \\ \mathbf{H}_I[n]: & \text{otherwise.} \end{cases}$$

The new set of hypotheses may be thought of as the first derivative of the previous hypotheses. Knowing the location of these events is equivalent to knowing the on/off trajectory of the primary signal. It turns out that this "event-based sensing" perspective can significantly improve the robustness of detection even for very sensitive detectors. After first giving the high-level intuition, a heuristic analysis is given in this section.

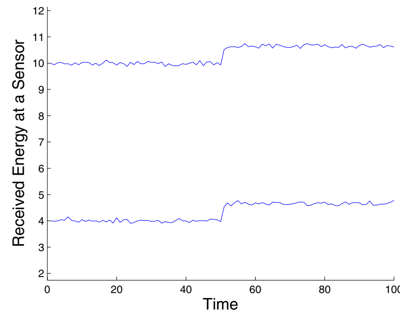

Fig. 4. Two sample-paths of noisy signal observations from the first noise model described in the "High Level Intuition" Section. The primary turns on in the middle of the plot. The traditional radiometer is very sensitive to uncertainty about the noise level but it is clear that event-based detection would work.

A. High-Level Intuition

We start with the simplistic scenario in which the noise average power is unknown but constant over time. Also assume that whenever the primary enters/exits, the received power changes instantaneously.[4] Furthermore, assume that in the time-scale of interest, the primary changes its state at most once. If the signal power is less than the noise uncertainty, the radiometer fails. On the other hand, no matter how weak the primary signal is, an edge detection algorithm will be able to sense the primary if the observation window is long enough (see Figure 4). This is exactly the time-domain analogy of the observation in [2] that noise-calibration leads to an elimination of the SNR wall whenever the noise is known to be perfectly white across frequencies.

Next, consider the more realistic scenario in which the noise is comprised of different interferers each of which have constant power over time but are allowed to turn on/off.[5] Furthermore, assume that whenever an interferer enters or exits, the power level changes instantaneously. Under this model, any single-user edge detector will also fail since it cannot distinguish between the entrance/exit of the primary and that of an interferer. However, notice

[3]When a secondary system is first turned on, it encounters a start-up transient within which the natural question is whether it is in a spatial spectrum hole at this time [11]. If it is impossible to answer this question because of an SNR wall, then simply being on for a very long time is not enough for a system to reach steady-state. One possibility is to use cooperation to bootstrap newly arrived systems into steady-state [11] and the other is to wait until we see two successive time-domain events: the arrival of the primary user followed by his departure. However, this second technique works only if we are certain that there is a maximum of one possible primary user.

[4]In reality, the detector's estimate of the received power is likely computed using a sliding window or some other filter that will smooth the transition somewhat.

[5]A physical example of this is the following scenario. The interferer is an 802.11 wireless router leaking very weak signals out of its assigned bandwidth because of imperfect filters. Because of the structure of packet-based transmissions, there will be an on/off pattern of interference in the band of interest.

that there *is* an improvement in the SNR wall. Instead of sensitivity being limited by the maximum number of possible interferers, the sensitivity is limited to the power corresponding to a *single* such interferer. This is because it is physically realistic to assume that unintended emitters are not going to be synchronized with each other — only one will turn on or off at any given time.

Even further improvements in sensitivity are possible with cooperation. The key observation is that in many realistic scenarios interferers are short-range whereas the primary signals have a large footprint[6] (call this the local-global assumption). The local-global assumption means that entrance/exit of the primary signal is sensed by most of the sensor nodes while any specific interferer's entrance/exit will only affect one or at most a few of the sensors at any given time. The physical effect is that of interference diversity. Therefore the system as a whole can detect the change in the activity status of the primary user whenever many sensors simultaneously detect it. We prove that as long as we do not deviate significantly from the local-global assumption, the proposed algorithm is able to correctly detect weak primary signals.

B. Formulation and Analysis

For a more rigorous formulation and analysis of the system see Appendix B. Here, an intuitive perspective is taken. Consider a network of sensors geographically distributed in an area. Assume that the sensors are connected to a Fusion Center (FC) through communication links that have limited capacity. Each sensor node transmits its belief about a possible change in the status of the primary transmitter (i.e. whether a primary has entered (or exited) or that no change has occurred). At each time instant, the FC will decide whether the primary has changed its state by, roughly speaking, taking a super-majority vote of the received messages from the sensors that vote for some change. For instance, if the fraction of sensor nodes believing that the primary has entered exceeds those believing that the primary has exited by more than a threshold, γ_{cu}, the FC will declare that the primary has entered.

The objective of an individual sensor is to detect the changes in the level of received power. It does this by comparing the measured power levels in the first and second halves of a sliding observation window. The goal of each sensor node is to provide the FC with as much information about the changes in the power level as

[6]This assumption does not hold when the primary users are themselves low-powered systems like wireless microphones. That case has many challenges [18].

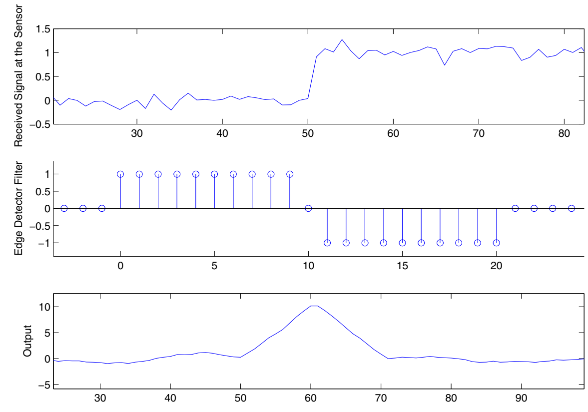

Fig. 5. The output of the edge-detector filter will be approximately zero unless there is a surge or drop in the received power level.

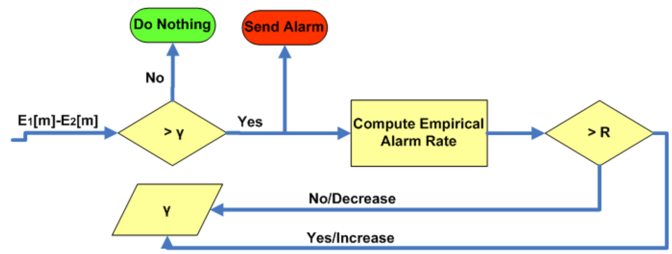

Fig. 6. Adaptive adjustment of the threshold at the sensor, γ is the threshold, R is the allowed data rate

possible. In reality, because of limited resources, every sensor is forced to send only the important changes by choosing some threshold γ to distinguish insignificant changes from significant ones. Figure 5 depicts a sample response of this edge detector to the entrance of a primary signal.

Figure 6 provides an adaptive approach for setting γ at each sensor node to avoid exceeding the control-channel rate limitation. Figure 7 shows a simulation illustrating the tradeoff between γ and the probability of false alarm under a specific noise model for different values of γ_{cu}.

Assume that the network consists of a large number of sensors and that the capacity of each link between a sensor node and the FC is some constant R. The local-global assumption (made precise in Appendix B) imposes certain constraints on the noise/interference signals observed at different nodes. It is important to note that the noise signals are *not* assumed to have a statistical description and can vary arbitrarily in time. The constraints imposed are therefore not on the statistical distributions of the noise signals but on their aggregate

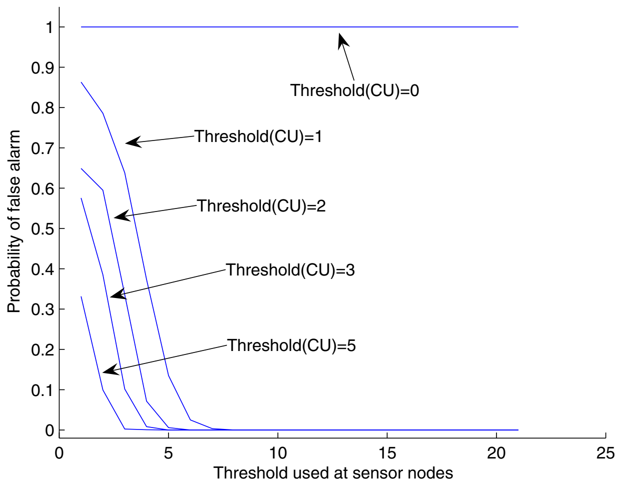

Fig. 7. The tradeoff between threshold at the sensors, threshold at the FC and probability of false alarm

average behavior in time or in space.

In a large network, we make the plausible assumption that the chance of an interferer exiting the system at some node of the network is the same as the chance of an interferer entering the network.[7] The local-global assumption states that the interferers are short-range. This means that the entrance or exit of an interferer affects only a small number of sensor nodes. As long as our assumptions hold, one can expect that when the on/off status of the primary user does not change, the number of sensors experiencing a surge in their received power level should be roughly equal to the number of sensors experiencing a drop in their received power level. This intuition is formally captured by the second constraint on noise in the Appendix's Definition 4.

The local-global assumption not only states that the interferers are local, but also that the primaries are global. This means that a change in the status of a primary is heard by many sensors. Roughly speaking, when the primary signal is not very weak, its entrance should increase the number of the sensors who experience a rapid increase in their received power level. Similarly, a primary leaving the system should increase the number of sensors who experience a rapid decrease in their received power level. This intuition is formally expressed as the second constraint on the noise in Definition 5.

Finally, the first constraint in Definition 4 implies that when the primaries do not enter or exit, the chosen threshold is sufficiently high so that the average number of alarms generated over time by each sensor node is bounded.

These definitions are not unrealistic. Both of the below scenarios fit[8] the model.

- When the noises (including interference noise) at different nodes are white, independently and identically distributed (i.i.d.) according to an unknown but continuous density function (such as a Gaussian).
- When the noise signal at each node has two components: 1. Independent white noise at different sensor nodes, identically distributed according to an unknown continuous density function; 2. Independent Poisson processes at different nodes with an unknown and possibly varying, but bounded rate modeling the entrance/exit of the interferers, each of which stays on for a fixed or iid random amount of time. While the interferer is on, it contributes an additional amount of white noise to the received signal.

If the power of the primary is very low, its arrival will not be able to excite enough sensors and hence the detection algorithm will fail. If the average power of the primary signal at the sensor nodes is above a certain limit, the algorithm should however succeed. Theorem 2 and its corollary formalizes this intuition.

The lower the power of the primary signal, the more sensitive the sensors need to be in detecting the changes in the power level. But increasing the sensitivity of the sensor nodes makes them vulnerable to the entrance/exit of weaker interferers. This would not only increase the number of false alarms generated by each sensor node (and thus create difficulties in communicating them all to the FC) but more importantly *gradually invalidates the local-global assumption*. A strong interferer may be heard by a larger group of sensor nodes. It should be noted however that the detection algorithm is robust as long as the local-global assumption remains valid. For a large network, the second constraint of Definition 4 remains valid as long as each interferer affects only a *sub-linear* number of sensors i.e. the fraction of affected sensors is relatively small as compared to the total number of sensors.

[7] This assumption does not always hold absolutely for the time-scales of interest. For example, in the morning, devices are more likely to be turned on and thus unintentional emitters are more likely to enter than they are to exit. However, if the emitters are themselves packet-oriented devices with short packets, then the packets will stop just as often as they start.

[8] The "fit" is in the sense that sample-paths of these random models will satisfy the definitions with very high probability.

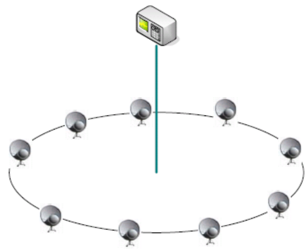

Fig. 8. In this example, the fusion center is placed equidistant from all sensor nodes.

"Computation in the air" method: Above, we assumed $R_1 = R_2 = ... = R_{m_s} = R$ allowing the total rate of control information $R_1 + R_2 + ... + R_{m_s}$ to approach infinity. This is not a realistic assumption. However, notice that the FC only needs to know the total number of positive and the total number of negative alarms rather than the precise identity of which nodes are sending positive or negative alarms. This echoes the setting of distributed computation combined with wireless communication (see for example [19]). It can be shown that separation between the communication stage and the computation stage is not always the optimal strategy, and a combined communication-computation scheme can generically improve the performance. This prompts the following strategy:

Consider the control channel (different from the channel being sensed) as a wireless link. The sensor nodes divide their control-channel resources into consecutive time intervals. In odd intervals, those nodes desiring to transmit a positive alarm simultaneously transmit signals of power P while the even intervals are used for transmitting negative alarm signals. The Fusion Center FC measures the power of the received signal (which is a superposition of the transmitted signals) in each of the intervals. The difference in the power levels of consecutive odd and even intervals provides an estimate of the difference between the number of positive and negative alarms.

Consider the simplistic model[9] shown in Figure 8 where the FC is equidistant from all of sensors and is experiencing the same channel to all of them on average. This makes sure that the sensor nodes are contributing equally on average to the total power received at the FC. If an α fraction of the m_s nodes are transmitting signals at power P, the expected value of the received power at the FC equals $\alpha \cdot P \cdot G_{av} + P_n$ where G_{av} is the average gain from the sensor nodes to the FC and P_n is the average power of the noise. Therefore the FC can estimate α based on the average received power. In the general scenario when the FC is placed on the ground, the channels between the FC and different sensor nodes would be different. The FC would receive a weighted average of the transmitted signals from different sensors. If the sensor nodes have some knowledge about the strength of their channel to the FC, they can adjust their power accordingly. Otherwise a more detailed investigation is required.

[9]The result can be probably extended to the following model: assume that we divide the sensor nodes into the set of green nodes and red nodes. When the green nodes are talking, the red nodes are listening. The red nodes estimate the number of green nodes who believe that the primary has entered/exited and thus play the role of the FC themselves.

IV. Simulations

The simulations are based on the following model: we consider a set of sensors uniformly distributed inside a square. At every point of the square lies an interferer that changes its status (i.e. turns on/turns off) with a fixed small probability. Around each sensor, we consider a smaller square. In this model, the interferers falling inside this square are sensed by the sensor without any error.

Figure 7 shows the tradeoff between the threshold at the sensors, the threshold at the FC and the probability of false alarm. As expected, a decrease in the threshold increases the number of alarms generated by each of the sensors when the signal is not present, and thus a higher chance of the FC giving a false alarm.

Figure 9 depicts the trade off between the number of sensors and the probability of false alarm. The goal of this simulation is twofold. Firstly, as long as the sensitivity area of sensors do not overlap, the alarms sent to the FC are zero mean i.i.d. random variables. As we increase the number of sensors, their average will be within an interval around zero with an increasing probability. Therefore the probability of false alarm will be reduced. Secondly, we observe that increasing the number of sensors in a closed area will eventually make them sensitive to common interferers. In other words, in a dense network, the alarms sent by the sensors will become correlated and the performance would not be improved by the same trend.

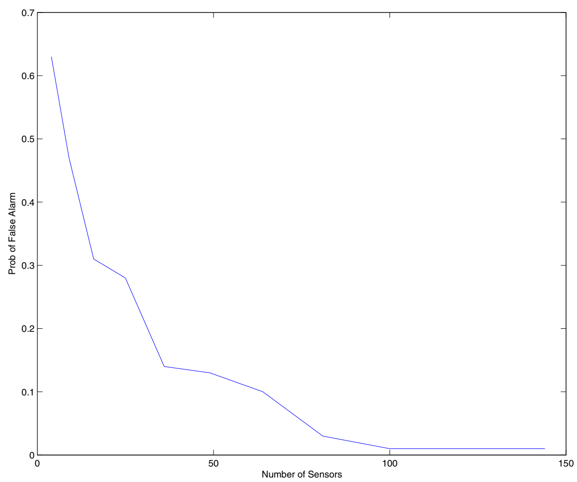

Fig. 9. Trade-off between number of sensors and probability of false alarm

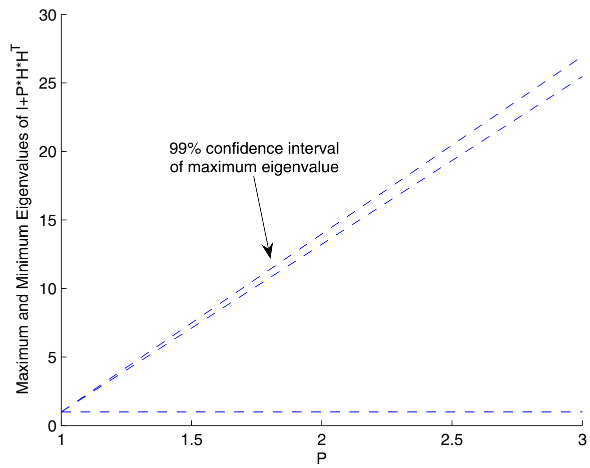

Fig. 10. Maximum and minimum eigenvalues of $I + pHH^T$ as a function of p, power of the signal. The minimum eigenvalue remains around one while maximum eigenvalue increases linearly. Here H is a random matrix. At $p = 0$, both eigenvalues are one. The maximum eigenvalue increases rapidly with the increase in p, while the minimum eigenvalue remains almost the same.

Appendix A
SNR Wall Analysis for the Max-Min Eigenvalue Detector

In this appendix, we prove that the possibility of external interferers can impose limits on the performance of the max-min eigenvalue detector [7]. In this section, we adopt the same notation used in [7] for reader convenience. The authors implicitly assume that the coherence time of the multi-antenna receiver is infinity. We continue with the same assumption.

Due to the existence of interferers even when the primary signal is not present, the ratio $\frac{\lambda_{\max}(K_{\widehat{\mathbf{X}}})}{\lambda_{\min}(K_{\widehat{\mathbf{X}}})}$ cannot be equal to one. This means that the detector has to set a threshold γ and decide between the hypothesis H_0 and the hypothesis H_1 based on $\mathbf{1}[\frac{\lambda_{\max}(K_{\widehat{\mathbf{X}}})}{\lambda_{\min}(K_{\widehat{\mathbf{X}}})} > \gamma]$. The value of γ would be constrained by the target probability of *false-alarm.*

Definition 1: Assuming any specific model for noise, the threshold function "Thf" is the answer to the following equation: $Pr\left(\frac{\lambda_{\max}(K_{\widehat{\mathbf{X}}})}{\lambda_{\min}(K_{\widehat{\mathbf{X}}})} > \text{Thf}(Pr_{FA})\right) = Pr_{FA}$.

Theorem 1: For any positive ϵ, there exists a threshold ρ such that whenever the primary power P is less than ρ, the detector cannot guarantee a probability of miss-detection less than $1-\epsilon$. Furthermore, given the probability of false alarm Pr_{FA}, ρ is bounded from below by $\frac{(\gamma-1)\sigma_\eta^2}{c_1(\gamma+1)}$ where $c_1 = O(\frac{1}{\epsilon})$ and $\gamma = Thf(Pr_{FA})$.

Proof: We prove the statement by contradiction. Fix some P_{FA} and some $\epsilon > 0$ and assume that the guarantee of $P_{MD} < 1-\epsilon$ exists. In particular, the sensor should have $P_{MD} < 1-\epsilon$ when the noise on the channels are white and uniform, and that $K_{\widehat{\mathbf{S}}} = P\mathbf{I}$. In this case, $K_{\widehat{\mathbf{X}}} = \kappa PHH^\dagger + \sigma_\eta^2\mathbf{I}_{ML}$ for some constant κ. Please note that intuitively if P is small here, from the perspective of the detector it behaves like an interferer and thus the ratio between maximum and minimum eigenvalue should be close to one.

Using the fact that for any matrices A and B we have:

$$\begin{cases} \lambda_{\max}(A+B) \leq \lambda_{\max}(A) + \lambda_{\max}(B); \\ \lambda_{\min}(A-B) \geq \lambda_{\min}(A) - \lambda_{\max}(B), \end{cases}$$

we can bound $\frac{\lambda_{\max}(K_{\widehat{\mathbf{X}}})}{\lambda_{\min}(K_{\widehat{\mathbf{X}}})}$ from above by $\frac{\sigma_\eta^2 + \kappa P\lambda_{\max}(HH^\dagger)}{\sigma_\eta^2 - \kappa P\lambda_{\max}(HH^\dagger)}$ as long as $\sigma_\eta^2 - \kappa P\lambda_{\max}(HH^\dagger) > 0$.

In order to use the Markov inequality in the above expression, we first need to show that $\mathrm{E}[\lambda_{\max}(HH^\dagger)] < \infty$.

Using Jensen's inequality we have:

$$\mathrm{E}[\lambda_{\max}(HH^\dagger)] \leq \mathrm{E}\sqrt{\sum_{i,j}|HH^\dagger_{i,j}|^2} \leq$$

$$\sqrt{\mathrm{E}\sum_{i,j}|HH^\dagger_{i,j}|^2} < \infty.$$

Now, using the Markov inequality, $Pr(\lambda_{\max}(HH^\dagger) \leq \frac{1}{\epsilon}\mathrm{E}\lambda_{\max}(HH^\dagger)) \geq 1-\epsilon$.

Let $c_1 = \kappa\frac{1}{\epsilon}\mathrm{E}\lambda_{\max}(HH^\dagger)$. The above inequality implies that if $P < \frac{\sigma_\eta^2}{c_1}$:

$$Pr\Bigg(\sigma_\eta^2 - \kappa P\lambda_{\max}(HH^\dagger) > 0 \text{ and}$$
$$\frac{\sigma_\eta^2 + \kappa P\lambda_{\max}(HH^\dagger)}{\sigma_\eta^2 - \kappa P\lambda_{\max}(HH^\dagger)} \leq \frac{\sigma_\eta^2 + c_1 P}{\sigma_\eta^2 - c_1 P}\Bigg) \geq 1-\epsilon.$$

Now, if $P < \frac{(\gamma-1)\sigma_\eta^2}{c_1(\gamma+1)}$, one can conclude that $\frac{\sigma_\eta^2+c_1P}{\sigma_\eta^2-c_1P} \leq \gamma$. This would in turn imply that

$$Pr(\frac{\sigma_\eta^2 + \kappa P\lambda_{\max}(HH^\dagger)}{\sigma_\eta^2 - \kappa P\lambda_{\max}(HH^\dagger)} \leq \gamma) \geq 1-\epsilon,$$

and therefore

$$Pr_{MD} = Pr(\frac{\lambda_{\max}(K_{\widehat{\mathbf{x}}})}{\lambda_{\min}(K_{\widehat{\mathbf{x}}})} \leq \gamma) \geq 1-\epsilon,$$

which is a contradiction. ■

APPENDIX B
FORMULATION AND ANALYSIS OF THE PROPOSED DETECTION ALGORITHM

A. Model Configuration

Consider a network of m_s sensors and one Fusion Center (FC). The sensors are geographically distributed in an area. Furthermore each sensor has m_a antennas. There are communication links between each sensor and the FC with rate constraints $R_1, R_2, ..., R_{m_s}$. The message set at each node is {positive alarm, negative alarm, no-change}. A positive (or negative) alarm corresponds to the sensor's belief that the primary has entered (or exited) while a no-change message corresponds to the belief that primary has not changed its state. At each time instant, the FC will decide whether the primary has changed its state based on the received messages from the sensors. In addition, there is a latency constraint D at the FC meaning that the system needs to detect state transitions of the primary within delay D. Let T be the minimum coherence time of all wireless communication channels. Assume that the minimum inter-transition duration of the primary signal is lower bounded by T_s, i.e. time between consecutive transitions of the primary is at least T_s.

We assume that there are at most m_t primaries. The active primaries are assumed to generate uncorrelated signals at any particular time instance. For simplicity, we also assume that all sensor nodes are synchronized in time. It should be noted that in a realistic scenario there will inevitably be some delay for reasons such as the coding used by the sensor nodes, possible clock mismatches etc.

TABLE III
NOTATIONS

Variable	Description
FC	Fusion Center
m_s	Number of sensor nodes
m_a	Number of antennas at each node
m_t	Number of primaries
R_i $(i = 1, 2, ..., m_s)$	Rate constraints
D	Latency constraint at the FC
T_s	lower bound on the minimum inter-trans. duration of primaries
T	minimum coherence times
γ_j	Threshold at the j^{th} sensor
γ_{cu}	Threshold at the FC
$\Phi_j(.)[m]$	Output of the j^{th} sensor
Dt(., ., .)[m]	Output of the FC

B. Detection Algorithm at the Sensors

In this algorithm, the objective of a sensor is to detect the changes in the received power due to entrance/exit of the primary. One possible scheme is to compare the measured power levels in the first and second halves of a shifting observation window. The length of the observation windows is chosen to be less than the minimum coherence time, T_c, ensuring that the changes in the power levels are not coming from changes in the channel statistics. Other constraints on the duration of the observation window are the network's latency constraint, D, and the minimum inter-transition duration, T_s. The latter constraint guarantees that the primary signal is active during the second half of some observation window.

Definition 2: For any positive γ, the output of the detection algorithm at the j^{th} sensor $\Phi_j(\gamma)$ is defined as follows: assuming that the received signal at the i^{th} antenna of the j^{th} sensor $(1 \leq i \leq m_a)$ is $x_{i,j}[.]$, the sensor transmits $\Phi_j(\gamma)[m] =$

$$\begin{cases} \text{Positive alarm} & \text{if } \mathbf{E}_{1,j}[m] - \mathbf{E}_{2,j}[m] > \gamma \\ \text{Negative alarm} & \text{if } \mathbf{E}_{1,j}[m] - \mathbf{E}_{2,j}[m] < -\gamma \\ \text{No-Change message} & \text{otherwise} \end{cases}$$

where

$\mathbf{E}_{1,j}[m] = \frac{1}{m_a}\sum_{r=1}^{m_a}\mathbf{E}_{2,j,r}[m] - \frac{1}{m_a}\sum_{r=1}^{m_a}\mathbf{E}_{1,j,r}[m]$,
$\mathbf{E}_{2,j}[m] = \frac{1}{m_a}\sum_{r=1}^{m_a}\mathbf{E}_{1,j,r}[m] - \frac{1}{m_a}\sum_{r=1}^{m_a}\mathbf{E}_{2,j,r}[m]$,
and
$\mathbf{E}_{1,j,r}[m] = \frac{1}{N_0}\sum_{n=0}^{N_0-1}|x_{r,j}[m-n-N_0-1]|^2$,
$\mathbf{E}_{2,j,r}[m] = \frac{1}{N_0}\sum_{n=0}^{N_0-1}|x_{r,j}[m-n]|^2$,
$2N_0 + 1 = \min(T_s, T_c, D)$. ●

We will discuss more about the choice of γ in a noisy environment later.

C. Detection Algorithm at the Fusion Center

At any time instance m, the Fusion Center receives signals from the sensor nodes. It seems intuitive for the FC to take a super-majority vote:

Definition 3: For any positive γ_{cu}, the output of the detection algorithm at the FC is defined as follows:

$$\mathbf{Dt}(\overrightarrow{\gamma}, \gamma_{cu}, \overrightarrow{R})[m] = \begin{cases} \mathbf{H}_+ : \\ \frac{1}{m_s}\sum_{j=1}^{m_s} \mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Positive alarm}] - \\ \frac{1}{m_s}\sum_{j=1}^{m_s} \mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Negative alarm}] > \gamma_{cu} \\ \mathbf{H}_- : \\ \frac{1}{m_s}\sum_{j=1}^{m_s} \mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Positive alarm}] - \\ \frac{1}{m_s}\sum_{j=1}^{m_s} \mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Negative alarm}] < -\gamma_{cu} \\ \mathbf{H}_I : \text{Otherwise} \end{cases}$$

where $\overrightarrow{\gamma} = (\gamma_1, \gamma_2, ..., \gamma_{m_s})$, $\overrightarrow{R} = (R_1, R_2, ..., R_{m_s})$.

D. Analysis

For simplicity of analysis, we assume that the channel is flat fading and that the primaries are using the phase-shift keying (PSK) digital modulation, i.e. the constellation points chosen are positioned with uniform angular spacing around the unit circle.

The received signal at the i^{th} antenna of the j^{th} sensor is equal to $x_{i,j}[n] = \sum_{z=1}^{m_t} h_{ij}^z s_z[n] + w_{i,j}[n]$ where h_{ij}^z is the channel gain and $s_z[n]$ is the signal transmitted by the z^{th} primary. If the primary is not active, $s_z[n] = 0$ otherwise it has constant amplitude $|s_z[n]|^2 = P_z$.

We have:

$$|x_{i,j}[n]|^2 = \sum_{z=1}^{m_t}\sum_{z'=1}^{m_t} h_{ij}^z s_z[n](h_{ij}^{z'} s_{z'}[n])^\dagger +$$

$$\sum_{z=1}^{m_t} h_{ij}^z s_z[n] w_{i,j}[n]^\dagger + \sum_{z=1}^{m_t} (h_{ij}^z s_z[n])^\dagger w_{i,j}[n] + |w_{i,j}[n]|^2$$

and so let $E_{*,j}[n] := \frac{1}{m_a}\sum_{i=1}^{m_a} |x_{i,j}[n]|^2$.

Averaging over different antennas cancels the effect of multipath fading. Since the transmitted signals from different primaries at a time instance are uncorrelated, using the law of large numbers we get the following equation: $\lim_{m_a\to\infty} E_{*,j}[n] = \sum_{z=1}^{m_t} \rho_{j,z}[n] P_z \mathbf{1}[z^{th} \text{ active at } n] + |W_{*,j}[n]|^2$ where $\rho_{j,z}[n]$ is the average gain from the z^{th} transmitter to j^{th} sensor node and $|W_{*,j}[n]|^2$ is the average noise power at the different antennas at time instance n. i.e. $|W_{*,j}[n]|^2 = \lim_{m_a\to\infty} \frac{1}{m_a}\sum_{i=1}^{m_a} |w_{i,j}[n]|^2$.

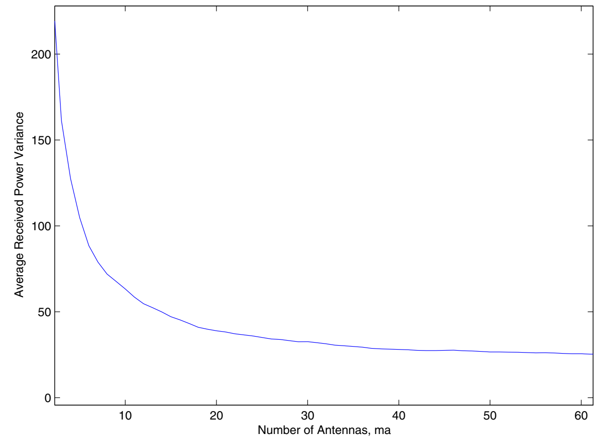

Fig. 11. Figure depicting the variance of $\frac{1}{m_a}\sum_{i=1}^{m_a} |w_{i,j}[n]|^2$ as a function of m_a. Here it is assumed that the noise signals is comprised of independent white Gaussian noises over different antennas, and also interference signals coming from an external interferer. As $m_a \to \infty$, the randomness coming from the white Gaussian noises and the fadings from the external interferer to the receiver decreases while the uncertainty in the power level of the samples of the external interferer remains.

Figure 11 depicts the average variance of $\frac{1}{m_a}\sum_{i=1}^{m_a} |w_{i,j}[n]|^2$ as a function of m_a for a specific noise model.

We therefore have: $\lim_{m_a\to\infty} \mathbf{E}_{1,j}[m] = \frac{1}{N_0}\sum_{n=0}^{N_0-1} E_{*,j}[m-n-N_0-1]$,
$\lim_{m_a\to\infty} \mathbf{E}_{2,j}[m] = \frac{1}{N_0}\sum_{n=0}^{N_0-1} E_{*,j}[m-n]$.

Observation 1: Since $2N_0 \leq T_c$, $\rho_{j,z}[n]$ does not change within an observation window. We can thus draw the conclusion that if the status of a primary does not change throughout the observation window, it does not contribute to the difference $\lim_{m_a\to\infty} \Phi_j(\gamma)[m] = \lim_{m_a\to\infty}(\mathbf{E}_{1,j}[m] - \mathbf{E}_{2,j}[m])$. ●

In the following subsections, we analyze the detection algorithm under different configurations. In the first subsection, we analyze the asymptotic behavior of the algorithm under the assumption that $R_1 = R_2 = ... = R_{m_s} = R$. We will prove that the FC will successfully detect weak primaries when $m_a, m_s \to \infty$ as long as the noise follows a certain model. The problem with this configuration is that the total rate, $R_1 + R_2 + ... + R_{m_s}$ approaches infinity as $m_s \to \infty$.

1) Equal Rate At All Sensors: Assume that $R_1 = R_2 = ... = R_{m_s} = R$. We will define a class of noise functions under which the detection algorithm has arbitrarily low probabilities of false alarm and missed-detection as m_s and m_a converge to infinity even for arbitrarily weak primaries. The constraint imposed on the noise functions quantifies the local-global assumption according to which the noise signals observed at different sensors are independent.

Definition 4: For any integer N_0 and vectors of positive real numbers $\overrightarrow{\gamma}$ and $\overrightarrow{R}$, Noise model $\Theta(\overrightarrow{\gamma}, \overrightarrow{R}, N_0, m_s, m_a)$ is the set of noise functions $w_{i,j}[n]$ ($1 \leq i \leq m_a$, $1 \leq j \leq m_s$, $n = ..., -2, -1, 0, 1, 2, ...$) satisfying the following two constraints whenever primaries are absent, i.e. $x_{i,j}[.] = w_{i,j}[.]$:
1. $\forall \; j, m \; : \; \lim_{M \to \infty} \frac{1}{M} \sum_{m=1}^{M} \mathbf{1}[\Phi_j(\gamma_j)[m] \neq \text{no-change}] \leq R$.
2. $\forall \; m$: $\lim_{m_s \to \infty} \left\{ \frac{1}{m_s} \sum_{j=1}^{m_s} \mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Positive alarm}] - \frac{1}{m_s} \sum_{j=1}^{m_s} \mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Negative alarm}] \right\} = 0$.

Remark : The first constraint ensures that for the chosen threshold γ_j, the uncoded transmissions of the alarms generated by the noise signal when primaries are absent do not violate the rate constraints[10] Observation 1 implies that when some of the primaries are present but do not change their state, the same alarm rate would be generated. Clearly increasing γ_j enlarges the set of noise sequences that satisfy the first property. In the extreme case of $\gamma_j = \infty$, the constraint would be valid for all R_j. The second assumption is intuitively saying that noise is unbiased: as $m_s \to \infty$, the number of positive alarms and negative alarms generated by the noise converge together. This phenomenon would happen if the noise sequences are each unbiased and behave independently of each other. If the local-global assumption fails, the alarms generated by different nodes will no longer be independent of each other and this constraint may no longer be valid.

Definition 5: For any constant c, the noise model $\Upsilon(c, \overrightarrow{\gamma}, N_0, m_s, m_a)$ is the set of noise functions $w_{i,j}[n]$ ($1 \leq i \leq m_a$, $1 \leq j \leq m_s$, $n = ..., -2, -1, 0, 1, 2, ...$) satisfying the following two constraints. Whenever primaries are absent, i.e. $x_{i,j}[.] = w_{i,j}[.]$, either $(\delta_1 > 0)$ or $(\delta_2 > 0)$ where:

$$\delta_1 = \inf_{m} \lim_{m_s \to \infty} \inf_{\overrightarrow{\gamma'} : \gamma_j \geq \gamma'_j \text{ for } j=1,2,..m_s \text{ and } \frac{1}{m_s} |\overrightarrow{\gamma} - \overrightarrow{\gamma'}|_1 \geq c} \left\{ \frac{1}{m_s} \sum_{j=1}^{m_s} \Big[\mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Positive alarm}] - \mathbf{1}[\Phi_j(\gamma'_j)[m] = \text{Positive alarm}] \Big] \right\};$$

and

$$\delta_2 = \inf_{m} \lim_{m_s \to \infty} \inf_{\overrightarrow{\widehat{\gamma}} : \gamma_j \leq \widehat{\gamma}_j \text{ for } j=1,2,..m_s \text{ and } \frac{1}{m_s} |\overrightarrow{\gamma} - \overrightarrow{\widehat{\gamma}}|_1 \geq c} \left\{ \frac{1}{m_s} \sum_{j=1}^{m_s} \Big[\mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Negative alarm}] - \mathbf{1}[\Phi_j(\widehat{\gamma}_j)[m] = \text{Negative alarm}] \Big] \right\}.$$

Remark : Clearly for any j and any $\gamma'_j \leq \gamma_j$,

$$\mathbf{1}[\Phi_j(\gamma_j)[m] = \text{Positive alarm}] \geq \mathbf{1}[\Phi_j(\gamma'_j)[m] = \text{Positive alarm}].$$

The first condition ensures that in the limit case, whenever the average gap between γ'_j and γ_j is at least a constant c, the rate of positive alarm is sensitive to changes in γ. The second condition has a similar interpretation. The reason for introducing these two conditions is that intuitively when a primary enters the system, the chance of passing threshold γ, and thus of "Positive alarm", increases whereas the chance of passing $-\gamma$, and thus of "Negative alarm", decreases. This is because the effective threshold for "Positive alarm" from the perspective of noise decreases, whereas the effective threshold for "Negative alarm" from the perspective of noise increases.

Theorem 2: The detection algorithm $\mathbf{Dt}(\overrightarrow{\gamma}, \gamma_{cu}, \overrightarrow{R})$ succeeds with probability one as $m_a, m_s \to \infty$ if

- The detector is operating in an environment with a noise sequence from $\Theta(\overrightarrow{\gamma}, \overrightarrow{R}, N_0) \bigcap \Upsilon(c, \overrightarrow{\gamma}, N_0, m_s, m_a)$;
- The average power of the primary that has entered/exited the system over all antennas of the system is at least c;
- $\gamma_{cu} < \max(\delta_1, \delta_2)$ where δ_1 and δ_2 are defined in Definition 5.

Proof: Assume that no primary changes its status in the interval $[n - N_0 - 1, n + N_0 - 1]$. The second property of $\Theta(\overrightarrow{\gamma}, \overrightarrow{R}, N_0)$ ensures that $\lim_{m_s \to \infty} \mathbf{Dt}(\overrightarrow{\gamma}, \gamma_{cu}, \overrightarrow{R}) = \mathbf{H}_I$.

Next assume that the z^{th} primary turns on at time n. Since $2N_0 + 1$ is less than the minimum inter-transition duration of the primary signals, T_s, there can be no other

[10] Notice here that we think of the rate constraint in terms of the probability of seeing a change. One could of-course use the binary entropy function to turn this into a bit-rate, but this would just be misleading in the kind of causal decision-making context considered here.

primary arriving or leaving in the interval $[n - N_0 - 1, n + N_0 - 1]$. Thus, other primaries do not contribute to $\lim_{m_a \to \infty} \Phi_j(\gamma)[m]$ according to Observation 1. Since $2N_0+1$ is less than the coherence time, $\rho_{j,z}[n]$ is constant during the interval $[n - N_0 - 1, n + N_0 - 1]$; let $\rho_j := \rho_{j,z}[u]$ (for $u \in [n - N_0 - 1, n + N_0 - 1]$).

The best opportunity to detect the entrance of the z^{th} primary is at time $n + N_0 - 1$, when the very beginning of the second half of the observation window is the entrance time of the z^{th} primary. The average power in the second part of the window is increased by $\rho_j P_z$. Let $\gamma'_j = \gamma_j - \rho_j P_z$, $\widehat{\gamma}_j = \gamma_j + \rho_j P_z$. Since the average power of the primary that has entered/exited the system over all antennas of the system is at least c, we get $\frac{1}{m_s} \sum_{j=1}^{m_s} \rho_j P_z \geq c$. The condition imposed in $\Upsilon(c, \overrightarrow{\gamma}, N_0, m_s, m_a)$ ensures that the average number of positive alarms that the FC receives is greater than the average number of negative alarms by at least $\max(\delta_1, \delta_2)$. This means that the detector $\mathbf{Dt}(\overrightarrow{\gamma}, \gamma_{cu}, \overrightarrow{R})$ succeeds in detecting the primary. ■

Corollary: The detector will be able to detect the entrance/exit of a primary signal with average power p_{min} if the noise signals belong to the set:

$$\bigcap_{c \geq p_{min}} \Upsilon(c, \overrightarrow{\gamma}, N_0, m_s, m_a) \bigcap \Theta(\overrightarrow{\gamma}, \overrightarrow{R}, N_0).$$

Acknowledgements

The authors would like to thank TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422) and the following organizations: Cisco, ESCHER, HP, IBM, Intel, Microsoft, ORNL, Pirelli, Qualcomm, Sun, Symantec, Telecom Italia and United Technologies, for their support of this work. The research was also partially supported by Sumitomo Electric and NSF grants CNS-0403427, CCF-0729122, CCF-0500023, ANI-0326503, CCF-0635372, and CNS-0627161. The authors also thank Prof. David Tse and the students in the graduate Wireless Communications course at Berkeley, along with the other students in Wireless Foundations and the Berkeley Wireless Research Center for giving feedback on this research, particularly Rahul Tandra, Mubaraq Mishra and Kristen Ann Woyach.

References

[1] S. Shellhammer, "Numerical Spectrum Sensing Requirements," IEEE Std. 802.22-06/0088r0, Jun. 2006.

[2] R. Tandra and A. Sahai, "SNR Walls for Signal Detection," *IEEE Journal on Special Topics in Signal Processing*, vol. 2, pp. 4-17, Feb. 2008.

[3] A. Sahai, N. Hoven and R. Tandra, "Some fundamental limits on cognitive radio," *Proc. Allerton Conf. on Control, Communication, and Computation*, Monticello, IL. Oct. 2004.

[4] A. Sahai and D. Cabric, "Spectrum sensing: fundamental limits and practical challenges," Tutorial presented at the *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Baltimore, MD. Nov. 2005.

[5] W. A. Gardner, "Spectral correlation of modulated signals: part I — analog modulation," *IEEE Trans. Communications*, vol. 35, no. 6, pp. 584-595, 1987.

[6] D. Cabric, S. M. Mishra and R. W. Brodersen, "Implementation Issues in Spectrum Sensing for Cognitive Radios," *Proc. of the Thirty-Eighth Asilomar Conference on Signals, Systems and Computers*, vol.1, pp. 772-776, November 2004.

[7] Y. Zeng and Y.C. Liang "Maximum-Minimum Eigenvalue Detection for Cognitive Radio", *The 18th Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications. (PIMRC)*, Sep. 2007.

[8] S. M. Mishra, A. Sahai and R. W. Broderson, "Cooperative sensing among cognitive radios," *Proc. IEEE International Conference on Communications (ICC)*, Istanbul, Turkey, Jun. 2006.

[9] A. Ghasemi and E. S. Sousa, "Collaborative Spectrum Sensing for Opportunistic Access in Fading Environments," *IEEE DySPAN*, Baltimore, USA, pp. 131-136, Nov. 2005.

[10] M. Matsui, H. Shiba, K. Akabane, and K. Uehara, "A Novel Cooperative Sensing Technique for Cognitive Radio," *PIMRC*, Sep. 2007.

[11] R. Tandra, S.M. Mishra and A. Sahai, "What is a spectrum hole and what does it take to recognize one?" To appear in the *Proceedings of the IEEE*, Dec. 2008.

[12] V. V. Veeravalli, "Decentralized quickest change detection," *IEEE Trans. on Inform. Theory*, vol.47, no.4, May 2001.

[13] Y. Mei, "Information Bounds and Quickest Change Detection in Decentralized Decision Systems", *IEEE Trans. on Inform. Theory*, Vol. 51, no.7, July 2005.

[14] M. Basseville and I. Nikiforov, *Detection of Abrupt Changes: Theory and Applications*. Prentice-Hall, Englewood Cliffs, NJ, 1993.

[15] T. Kailath and H.V. Poor, "Detection of Stochastic Processes," *IEEE Trans. on Inform. Theory*, Vol. 44, pp. 2230-2259, Oct. 1998.

[16] H. Li, C. Li and H. Dai, "Quickest Spectrum Sensing in Cognitive Radio," *Proc. Conf. on Information Sciences and Systems (CISS)*, Mar. 2008.

[17] A. Sahai, N. Hoven, S.M. Mishra and R. Tandra, "Fundamental tradeoffs in robust spectrum sensing for opportunistic frequency reuse," Technical Report, Mar. 2006.

[18] S.M. Mishra, R. Tandra and A. Sahai, "Coexistence with primary users of different scales," *IEEE DySPAN*, Dublin, Ireland. pp. 158-167, Apr. 2007.

[19] B. Nazer and M. Gastpar, "Computation over Multiple-Access Channels," *IEEE Trans. on Inform. Theory*, vol.53, no.10, pp. 3498-3516, Oct. 2007.