Unified Space-Time Metrics to Evaluate Spectrum Sensing

Rahul Tandra, Qualcomm

Anant Sahai, University of California, Berkeley

Venugopal Veeravalli, University of Illinois at Urbana-Champaign

ABSTRACT

Frequency-agile radio systems need to decide which frequencies are safe to use. In the context of recycling spectrum that may already be in use by primary users, both the spatial dimension to the spectrum sensing problem and the role of wireless fading are critical. It turns out that the traditional hypothesis testing framework for evaluating sensing misses both of these and thereby gives misleading intuitions. A unified framework is presented here in which the natural ROC curve correctly captures the two features desired from a spectrum sensing system: safety to primary users and performance for the secondary users. It is the trade-off between these two that is fundamental. The spectrum holes being sensed also span both time and space. The single-radio energy detector is used to illustrate the tension between the performance in time and the performance in space for a fixed value of protection to the primary user.

INTRODUCTION

Philosophically, frequency-agile radios' spectrum sensing is a binary decision problem: is it safe to use a particular frequency where we are, or is it unsafe? So it seems natural to mathematically cast the problem as a binary hypothesis test. Most researchers model the two hypotheses as *primary user present* and *primary user absent*. This suggests that the key metrics should be the *probability of missed detection* P_{MD} and the *probability of false alarm* P_{FA}. But is this truly the right model? Does it illuminate the important underlying trade-offs?

To understand how metrics can matter, it is useful to step back and consider familiar capacity metrics. Traditionally, the community studied point-to-point links. There, Shannon capacity (measured in bits per second per Hertz) is clearly the important metric. However, this is not enough when we consider a wireless communication network — the spatially distributed aspect is critical, and this shows up in the right metrics. For instance, Alouini and Goldsmith in [1] propose the *area spectral efficiency* (measured in bits per second per square kilometer per Hertz) when links are closely packed together, and Gupta and Kumar in [2] further propose the *transport capacity* (measured in bit-meters per second per Hertz) when cooperation (e.g., multihop) is possible. It is these metrics that give much deeper insights into how wireless communication systems should be designed.

Spectrum sensing is about recycling bandwidth that has been allocated to primary users and thereby increasing the capacity pre-multiplier for the secondary system. There turns out to be a significant spatial dimension to spectrum recycling for a simple reason — the same frequency will be reused by another primary transmitter once we get far enough away. Thus, the potential *spectrum holes* span both time and space.

To see why ignoring this spatial dimension is misleading, we must first review the traditional binary hypothesis testing story where the central concept is *sensitivity*: the lowest received signal power for which target probabilities of false alarm and missed detection can be met. More sensitive detectors are considered better and it is well known that sensitivity can be improved by increasing the sensing duration.

However, why does one demand very sensitive detectors? The strength of the primary's signal is just a proxy to ensure that we are *far enough*. If wireless propagation were perfectly predictable, then there would be a single *right level* of sensitivity. It is the reality of fading that makes us demand additional sensitivity. But because fading can affect different detectors differently, a head-to-head comparison of the sensitivity of two detectors can be misleading. Instead, the possibility of fading should be incorporated into the *signal present* hypothesis itself.

The bigger conceptual challenge comes in trying to understand *false alarms*. The traditional hypothesis-test implicitly assumes that a false alarm can only occur when the primary users are entirely absent. But in the real world, the spectrum sensor must also guard against saying that it is close to the primary when it is far enough away. The *signal absent* hypothesis needs to be modified in some reasonable way that reflects both these kinds of false alarms. We must take into account that the users doing the sensing have some spatial distribution.

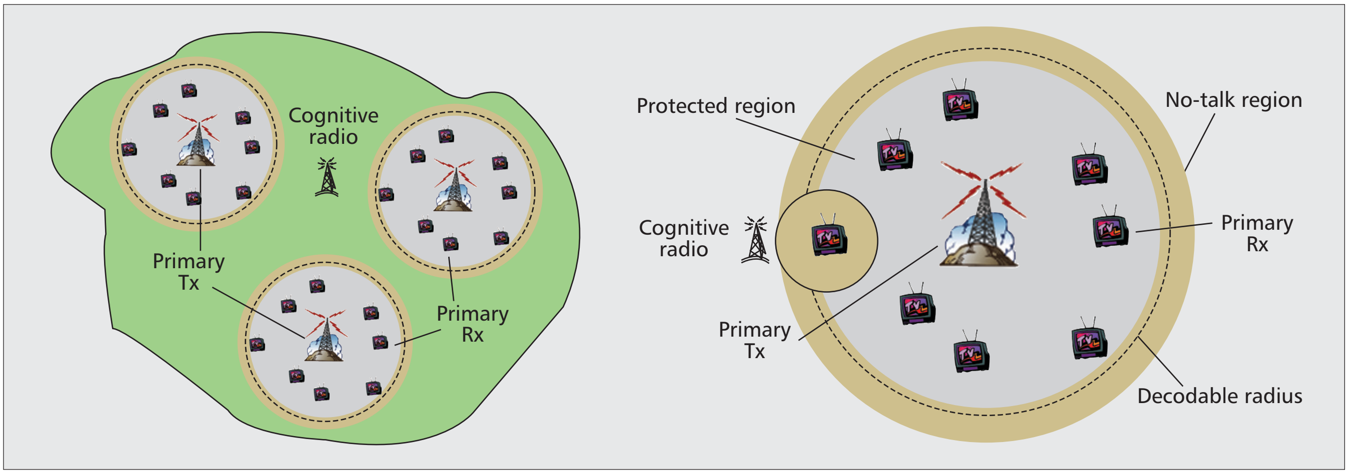

Figure 1. *This figure illustrates the scenario of cognitive radios acting as sensing-based secondary users recycling TV whitespaces. The secondary user is allowed to use the channel if it is outside both the protected region and the no-talk region (the tan-colored annulus shown in the figure) of each primary transmitter that is currently ON. The spectrum-sensing problem boils down to identifying whether the secondary user is within a space-time spectrum hole or not.*

Once both hypotheses have been appropriately modified, the receiver-operating-characteristic (ROC) curve appropriately reflects the fundamental trade-off in spectrum sensing between the safety guarantee for the primary users (captured by a metric we call the *Fear of Harmful Interference*, F_{HI}) and the secondary users' ability to recycle the leftover spectrum for themselves (captured by the *Weighted Probability of Space Time Recovered* metric, *WPSTR*). However, there are two subtle, but important, issues that must be addressed along the way lest we end up with trivial trade-offs. Both of these have to do with understanding the *nature* of the safety guarantee to the primary users. First, the underlying probabilistic model regarding the spatial distribution of the secondary users should not be consistent across the two hypotheses. In fact, it is better to use a worst-case spatial distribution under the *frequency band occupied* hypothesis so that the primary users' safety guarantee is strong. Second, the safety guarantee to primary users needs to be weakened at the start of each primary ON period. A time-domain sacrificial buffer zone needs to be introduced within which interference from secondary users is permissible; this gives the secondary user some time to evacuate the band and thus allows for some sensing. Without such a sacrificial buffer, the trade-offs almost invariably become trivial [3].

Unlike the traditional sensitivity-oriented metrics, these new metrics give a unified framework to compare different spectrum-sensing algorithms and yield several new insights into the space-time sensing problem. First, they clearly show that fading uncertainty forces the *WPSTR* performance of single-radio sensing algorithms to be very low for desirably small values of F_{HI}. This captures the fact that a single radio examining a single frequency cannot distinguish whether it is close to the primary user and severely shadowed, or if it is far away and not shadowed. Second, the metrics reveal the importance of diversity and how simple non-coherent detection can outperform matched filters in practice. Third, an example is used to show that there exists a non-trivial trade-off between the spatial and temporal performance for a spectrum sensor. In general, there exists an optimal choice of the sensing time for which the *WPSTR* metric is maximized.

Spectrum Sensing by Secondary Users

Spectrum holes [4], are regions in space, time and frequency within which it is safe for a secondary radio system to recycle the spectrum. The picture on the left in Fig. 1 shows there is a spatial region around every primary transmitter, called the no-talk region, within which secondary users are not allowed to transmit. The spectrum hole is everywhere else — shown here in green.

Intuitively, the two important dimensions along which a sensing algorithm should be evaluated are: the degree to which it is successful in identifying spectrum holes that are actually there; and the amount of harmful interference caused to the primary system by falsely identifying spectrum holes. An ideal approach — for example, involving a centralized database with primary user participation and geolocation functionality for secondary users [5] — would, by definition, create zero unauthorized harmful interference and yet recover all the spectrum holes.

To make the problem concrete, we now focus on a single primary user transmitting on a given frequency band. The picture on the right in Fig. 1 illustrates that the primary transmitter (a TV tower in this example) has a protection region (gray region in the figure), and any potential primary receivers within this area must be protected from harmful interference. The resulting no-talk radius r_n can be computed from the protection radius, the transmit power of the secondary radio, and the basic wireless propagation model [6]. The sensing problem thus boils down to deciding whether the distance from the TV tower is less or greater than r_n.

Review of the Traditional Time-Domain Formulation for Sensing

Currently, the most popular formulation of the spectrum sensing problem casts it as a binary hypothesis test between the following two hypotheses: *primary ON* and *primary OFF*.

The two traditional hypotheses are:

$$\text{Signal absent } \mathcal{H}_0 : Y[n] = W[n]$$
$$\text{Signal present } \mathcal{H}_1 : Y[n] = \sqrt{P}X[n] + W[n], \quad (1)$$

for $n = 1, 2, \cdots, N$. Here P is the received signal power, $X[n]$ are the unattenuated samples (normalized to have unit power) of the primary signal, $W[n]$ are noise samples, and $Y[n]$ are the received signal samples.

The two key metrics in this formulation are: the probability of false alarm, P_{FA}, which is the chance that a detector falsely thinks that the signal is present given that the signal is actually absent; and the probability of missed detection, P_{MD}, which is the chance that the detector incorrectly declares the signal to be absent given that the signal is actually present.

The lowest signal power P at which the detector can reliably meet (P_{FA}, P_{MD}) targets is called the detector's *sensitivity*. Furthermore, the minimum number of samples required to achieve a target sensitivity is called the detector's *sample complexity*. The traditional metrics triad of sensitivity, P_{FA}, and P_{MD}, are used along with the sample complexity to evaluate the performance of detection algorithms.

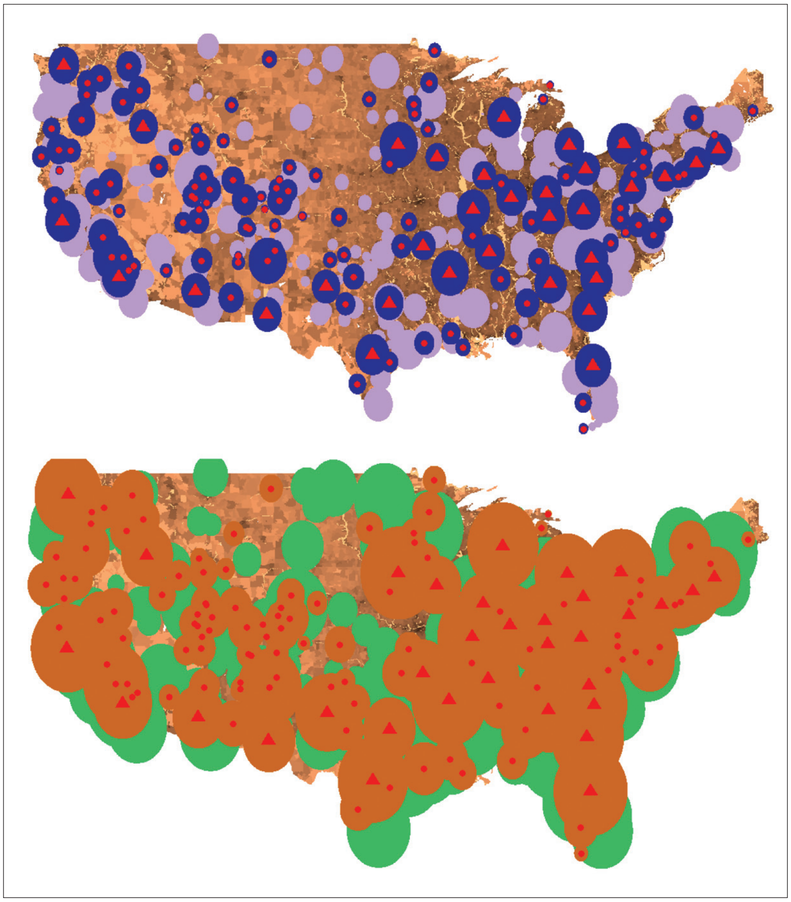

Figure 2. *The map on the top shows the location of TV towers (red triangles and circles) transmitting on channel 39 in the continental United States. The larger disks around the transmitters show the no-talk region around the TV transmitters within which a secondary user cannot recycle the channel. This shows that the true spectrum hole covers about 47 percent of the total area. The effective no-talk region for a radio using the –114 dBm rule (from [5]) is shown in the bottom figure — only 10 percent of the total area remains. This figure is taken from [7] where more details can be found on the available whitespace spectrum in TV bands.*

Drawbacks with the Traditional Formulation

The key idea behind the formulation in the previous section is that a detector that can sense weak signals will ensure an appropriately low probability of mis-declaring that we are outside the no-talk radius whenever we are actually inside. However, this formulation has some fundamental flaws.

How Much Sensitivity Do We Really Need? — The right level of sensitivity should correspond to the signal power at the no-talk radius. If there were no fading, the required sensitivity would immediately follow from the path-loss model. The traditional approach to deal with fading is to incorporate a fading margin into the target sensitivity (e.g., set the sensitivity low enough to account for all but the 1 percent worst case fades). However, different detectors may be affected differently by the details of the fading process. For example, a coherent detector looking for a single pilot tone would require a larger fading margin than a non-coherent detector averaging the signal power over a much wider band. Hence, thinking in terms of a single level of sensitivity for all detectors is flawed.

How to Measure the Performance of Spectrum Sensors — Traditionally, the *frequency unused* hypothesis ($\mathcal{H}_0$) has been modeled as receiving noise alone. However, it is perfectly fine for the primary transmitter to be ON, as long as the spectrum sensor can verify that it is outside the primary's no-talk radius. The real world hypothesis $\mathcal{H}_0$ is actually different at different potential locations of the secondary radio.

Building in a fading margin to the sensitivity has the unfortunate consequence of causing the false-alarm probabilities to shoot up when the spectrum sensor is close, but not too close, to the primary transmitter. This makes parts of the spectrum hole effectively unrecoverable [7]. Figure 2 shows this effect in the real world.

Spectrum Sensing: A Space-Time Perspective

The discussion in the previous section forces us to rethink the traditional hypothesis-testing formulation. Fading must be explicitly included and the reality of different potential locations must also be explicitly accounted for.

The received signal can be modeled as $Y[n] = \sqrt{P(R)}\, X[n] + W[n]$ whenever the primary transmitter is ON, where R is the distance of the secondary radio from the primary transmitter. The received signal *power* $P(R)$ is actually a random operator (modeled as independent of both the normalized transmitted signal and the noise) that depends on both the path loss and fading distributions. This gives the following composite hypothesis testing model:

$$\mathcal{H}_0 : Y[n] = \begin{cases} \sqrt{P(R)}X[n] + W[n] & R > r_n \text{ and primary ON} \\ W[n] & \text{primary OFF} \end{cases}$$
$$\mathcal{H}_1 : Y[n] = \sqrt{P(R)}X[n] + W[n] \quad R \in [0, r_n], \tag{2}$$

where we still need to decide on the primary user's ON/OFF behavior and the distributions for R in the two hypotheses to permit evaluation of the two kinds of error probabilities for any spectrum sensor.

Modeling Space

The true position of the secondary user relative to the primary transmitter is unknown. This is why we are sensing. For $\mathcal{H}_1$, it is natural to assume a worst-case position, generally at just within r_n where the primary signal is presumably weakest. A worst-case assumption makes the quality guarantee apply uniformly for all the protected primary receivers.

Suppose we took the same approach to $\mathcal{H}_0$. Typically, the worst case location under $\mathcal{H}_0$ would be just outside r_n with the primary user ON. After all, if we can recover this location, we can presumably recover all the locations even further away or when the primary user is OFF. Alas, this approach is fatally flawed since the signal strength distributions just within r_n and just outside of r_n are essentially the same. No interesting trade-off is possible because we are missing the fundamental fact that a sensing-based secondary user must usually give up some area immediately outside of r_n to be sure to avoid using areas within r_n.

Simply averaging over the distance R also poses a challenge. The interval (r_n, ∞) is infinite and hence there is no uniform distribution over it. This mathematical challenge corresponds to the physical fact that if we take a single primary-user model set in the Euclidean plane, the area outside r_n that can potentially be recovered is infinite. With an infinite opportunity, it does not matter how much of it we give up!

In reality, there are multiple primary transmitters using the same frequency. As a radio moves away from a given primary transmitter (R increases), its chance increases of falling within the no-talk radius of an adjacent primary transmitter. The picture on the bottom in Fig. 3 illustrates the Voronoi partitioning of a spatially distributed network of primary transmitters, and the picture on the top shows the effective single primary transmitter problem with a finite area.

The key is to choose a probability measure $w(r)rdr$ so as to weight/discount area outside r_n appropriately. The rigorous way to do this is to use results from stochastic geometry and point-process theory [8]. However, the key insights can be obtained by choosing any reasonable probability measure. The numerical results here have been computed assuming $w(r)$ is constant ($= c$) for $0 < r \leq r_n$, and an exponential weighting function, $w(r) = A\exp(-\kappa(r - r_n))$, for $r > r_n$. The constant part essentially tells us the probability of the primary being OFF. An exponential distribution is chosen for the rest because it has the *maximum entropy* among the set of all distributions on $[r_n, \infty)$ with a given mean. In our case, this mean is related to the average minimum distance between two primary towers transmitting on the same channel.

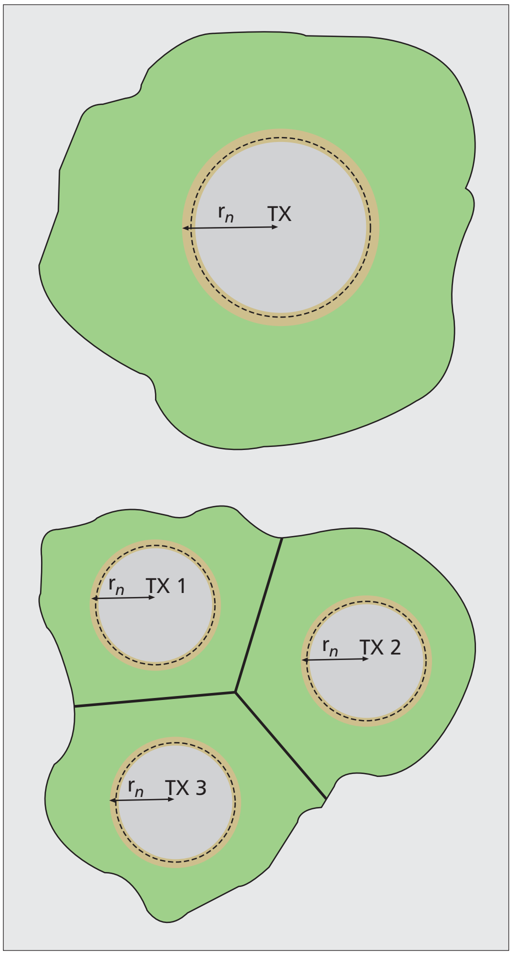

Figure 3. *The picture on the bottom shows the Voronoi partitioning of the space between primary transmitters. The multiple primary transmitter problem is approximated as an ideal single primary problem by including a spatial weighting function* w(r) *that discounts the value of areas far away from the primary transmitter.*

Building in a fading margin to the sensitivity has the unfortunate consequence of causing the false alarm probabilities to shoot up when the spectrum sensor is close, but not too close, to the primary transmitter. This makes parts of the spectrum hole effectively unrecoverable.

Modeling Time

The probability of being within the no-talk radius in $\mathcal{H}_0$ seems to capture the ON/OFF behavior in a long-term average sense. But long-term averages are not enough to allow us to evaluate sensing. Intuitively, if the primary user is coming and going very often, the issue of *timeliness* in sensing is more important than when the primary user is like a real television station and switches state very rarely, if at all.

Consider a secondary user that is located inside the no-talk radius. Let $U[n] = 1$ only if the primary transmitter is ON at time instance n. Assume that we start sensing at time instances n_i, and at the end of each sensing interval, the secondary user makes a decision of whether the frequency is safe to use ($D_i = 0$) or not safe to use ($D_i = 1$). The secondary user transmits only

The naive definition does not recognize that causality implies that the initial segments of a primary transmission are intrinsically more exposed to interference. This is the time-domain counterpart to the spatial status of primary receivers located at the edge of decodability.

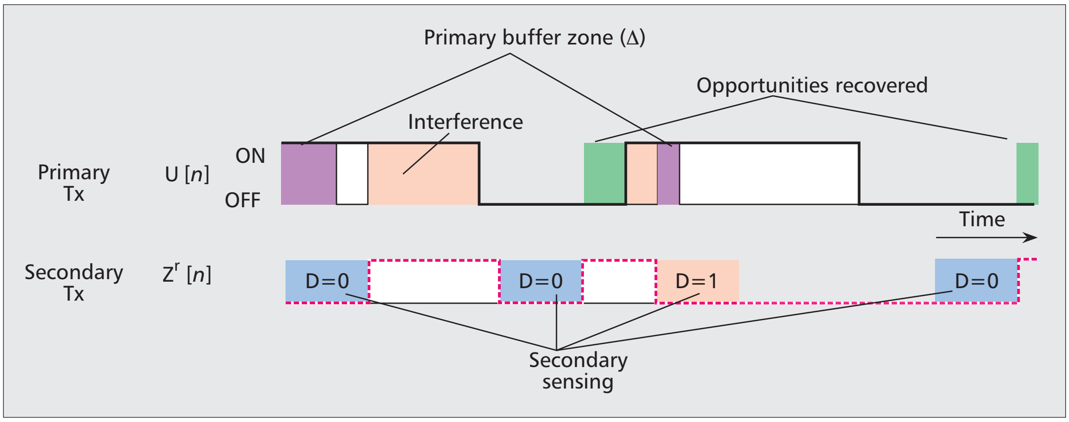

Figure 4. *The state of the primary user* U[n]*, the sensing epochs, as well as the secondary ON/OFF process* $Z^r[n]$ *(dashed red line) are shown in the figure. The red sensing windows indicate events when the detector declares the frequency to be used, and the blue sensing windows indicate when the detector declares the frequency to be unused. The primary sacrificial buffer zones are shown by the purple shaded regions on the function* U[n]*, and the actual harmful interference events are shown by shaded tan regions on* U[n].

if the frequency is deemed to be safe, and then senses again. This induces a random process $Z^r[n] \in \{0, 1\}$ denoting the state of a secondary user located at a distance r from the primary transmitter, with 1 representing an actively transmitting secondary user. An example scenario is shown in Fig. 4.

Intuitively, harmful interference could be quantified by measuring the fraction of the primary *ON* time during which a secondary user located inside the no-talk radius is transmitting. Suppose that the primary transmitter is OFF, a secondary user senses for N samples, correctly declares that the primary is OFF, and hence starts transmitting. There is still a finite non-zero probability that the primary comes back ON *while* the secondary is transmitting. This probability depends on the duration of the secondary user's transmission, but might have no connection to how long N is! For example, there would indeed be no connection in a Poisson model (the maximum-entropy modeling choice) for primary transmissions [3].

The secondary could thus cause interference even when its spectrum sensor is as correct as it could possibly be. If this definition of interference were to be adopted, the only way to drive the probability of interference to zero would be to scale the secondary transmission time to zero. This would give a relatively uninteresting trade-off between the protection to the primary system and the performance of the secondary user.

The naive definition does not recognize that causality implies that the initial segments of a primary transmission are intrinsically more exposed to interference. This is the time-domain counterpart to the spatial status of primary receivers located at the edge of decodability. Just as these marginal receivers must be sacrificed for there to be meaningful spectrum holes, it makes sense to assume that there is a temporal sacrificial buffer zone (Δ samples long) at the beginning of every OFF to ON transition of the primary user (illustrated as purple regions in Fig. 4). Secondary transmissions during this time should not be considered harmful interference.

SPACE-TIME METRICS

We now define two key metrics that are similar to the traditional metrics of P_{FA} and P_{MD}, but are computed on the composite hypotheses in Eq. 2. The trade-off between them is the fundamental *ROC* curve for the problem of spectrum sensing.

Safety: Controlling the Fear of Harmful Interference — This metric measures the worst-case safety that the sensing-based secondary user can guarantee to the primary user under uncertainty. We call it the *Fear of Harmful Interference* $F_{HI} = \sup_{0 \le r \le r_n} \sup_{F_r \in \mathbb{F}_r} \mathcal{P}_{F_r}(D = 0 | R = r)$, where $D = 0$ is the detector's decision declaring that the frequency is safe to use, and $\mathbb{F}_r$ is the set of possible distributions for $P(r)$ and $W[n]$ at a distance of r from the primary transmitter. The outer supremum is needed to issue a uniform guarantee to all protected primary receivers. The inner supremum reflects any non-probabilistic uncertainty in the distributions of the path-loss, fading, noise, or anything else.

Performance: Success in Recovering Spectrum Holes — By weighting the probability of finding a hole $P_{FH}(r)$ with the spatial density function, $w(r)r$, we compute the *weighted probability of space-time recovered* (*WPSTR*) metric:

$$WPSTR = \int_0^\infty P_{FH}(r) w(r) r \, dr, \text{ where}$$

$$P_{FH}(r) = \begin{cases} \lim_{M\to\infty} \frac{1}{M} \sum_{n=1}^{M} \mathbf{I}_{(Z^r[n]=1)} & \text{if } r > r_n \\ \lim_{M\to\infty} \frac{\sum_{n=1}^{M} \mathbf{I}_{(Z^r[n]=1, U[n]=0)}}{\sum_{n=1}^{M} \mathbf{I}_{(U[n]=0)}} & \text{if } r > r_n \end{cases}, \quad (3)$$

and $\int_0^\infty w(r) r dr = 1$. The $\mathbf{I}$ above is shorthand for indicator functions that take the value 1 whenever their subscript is true and 0 otherwise. Notice that the integral spans locations inside and outside the no-talk radius (0 to ∞). The name *WPSTR* reminds us of the weighting of performance over space and time. 1 – *WPSTR* is the appropriate analog of the traditional P_{FA}, except that it also implicitly includes the overhead due to the sample-complexity.

INSIGHTS FROM THE NEW SPACE-TIME FRAMEWORK

ALWAYS ON PRIMARIES: PURELY SPATIAL SPECTRUM HOLES

Assume that the primary transmitter is always *ON*. This corresponds to sensing a spectrum hole whose temporal extent is infinite, so it does not matter how long we spend sensing. To remind us of this spatial focus (and to maintain consistency with [4]), we call the secondary performance metric the *Weighted Probability of Area Recovered*, *WPAR*, instead of *WPSTR*.

Consider a single secondary user running a perfect radiometer (i.e., one with an infinite number of samples). If the noise variance is perfectly known, it is straightforward to derive expressions for F_{HI} and *WPAR* [4]. The black curve in Fig. 5 shows the F_{HI} vs. *WPAR* trade-off. Notice that the *WPAR* performance at low F_{HI} is bad even for the perfect radiometer. This captures the physical intuition that guaranteeing strong protection to the primary user forces the detector to budget for deep non-ergodic fading events. Unlike in traditional communication problems where there is no harm if the fading is not bad, here there is substantial harm since a valuable spectrum opportunity is left unexploited.

Impact of Noise Uncertainty: SNR Walls in Space — The real world noise power is not perfectly known. In the traditional formulation, this uncertainty causes the radiometer to hit an *SNR wall* that limits its sensitivity [9]. What happens under these new metrics?

See [10] for the details, but the result is illustrated by the red curve in Fig. 5. The noise uncertainty induces a critical F_{HI} threshold below which none of the spectrum hole can be recovered (*WPAR* = 0). In traditional terms, the sensitivity required to budget for such rare fading events is beyond the SNR wall. Just as the sample-complexity explodes to infinity as the SNR Wall is approached in terms of traditional sensitivity, the area recovered crashes to zero as F_{HI} approaches this critical value.

Dual Detection: How to Exploit Time-Diversity — The true power of these new metrics is that they allow us to see the importance of diversity. This can be cooperative diversity as discussed in [4], but the effect can be seen even with a single user. For example, one could presumably exploit time diversity for multipath if we believed that the actual coherence time is finite $N_c < \infty$. However, for the radiometer, all the thresholds must be set based on the primary user's fear of an infinite coherence time — the spectrum sensor might be stationary. The radiometer cannot do anything to exploit the likelihood of finite coherence times even if the sensor is likely to be moving.

The situation is different for a sinusoidal pilot tone, as illustrated by the blue curve in Fig. 5. The best-case scenario for coherent detection from a traditional sensitivity perspective — infinite coherence time with no noise uncertainty — can be *worse* in practice than a simple radiometer with noise uncertainty. As the sinusoidal pilot is narrowband, the matched filter suffers from a lack of frequency diversity as compared to the radiometer: fading is more variable and the resulting conservatism costs us area.

So does the matched filter have any use in wideband settings? Yes. It gives us an opportunity to deal with uncertain coherence-times. We can run two parallel matched filters — one assuming an infinite coherence time and the other doing non-coherent averaging to combine matched-filtered segments of length N_c — with their thresholds set according to their respective assumptions on the coherence time. If either of them declares that the frequency is used, then the secondary user will not use this frequency. This ensures that the F_{HI} constraint is met irrespective of the actual coherence time. The dual-detector approach thus leads to different F_{HI} vs. *WPAR* curves depending on what the mix of underlying coherence times is (stationary devices or moving devices).

The dashed curve in Fig. 5 shows the performance of the matched filter running with a known coherence time of N_c. Because it enjoys time-diversity that wipes out multipath fading, it is only limited by the same non-ergodic wideband shadowing that limits even the wideband non-coherent detector without noise-uncertainty. In principle however, this dual detector still has an SNR wall due to noise uncertainty. However, to be able to illustrate this effect, Fig. 5 was plotted using a very short coherence time N_c = 100. For any realistic coherence time, the *SNR* wall effect would become negligible at all but extremely paranoid values for F_{HI}.

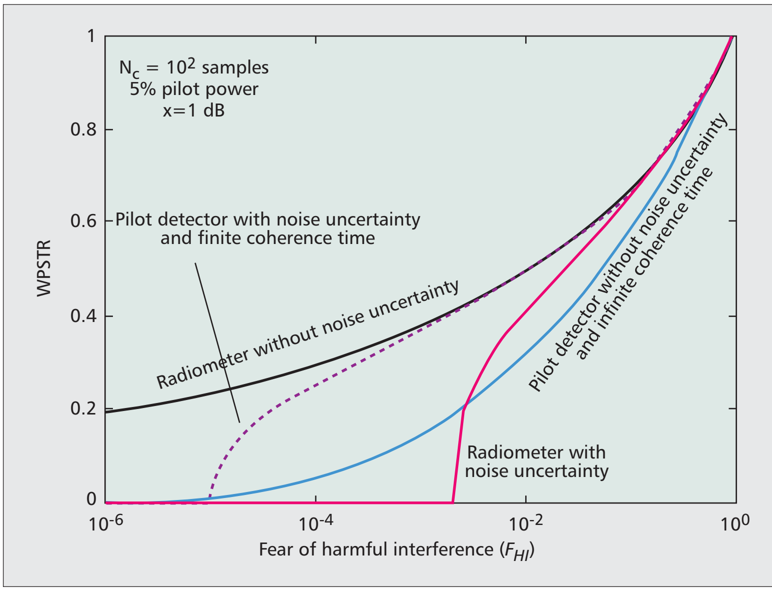

Figure 5. *Under noise uncertainty (1 dB here), there is a finite* F_{HI} *threshold below which the area recovered by a radiometer is zero* (WPSTR = *0). The coherent detector (modified matched filter) has a more interesting set of plots discussed in this article.*

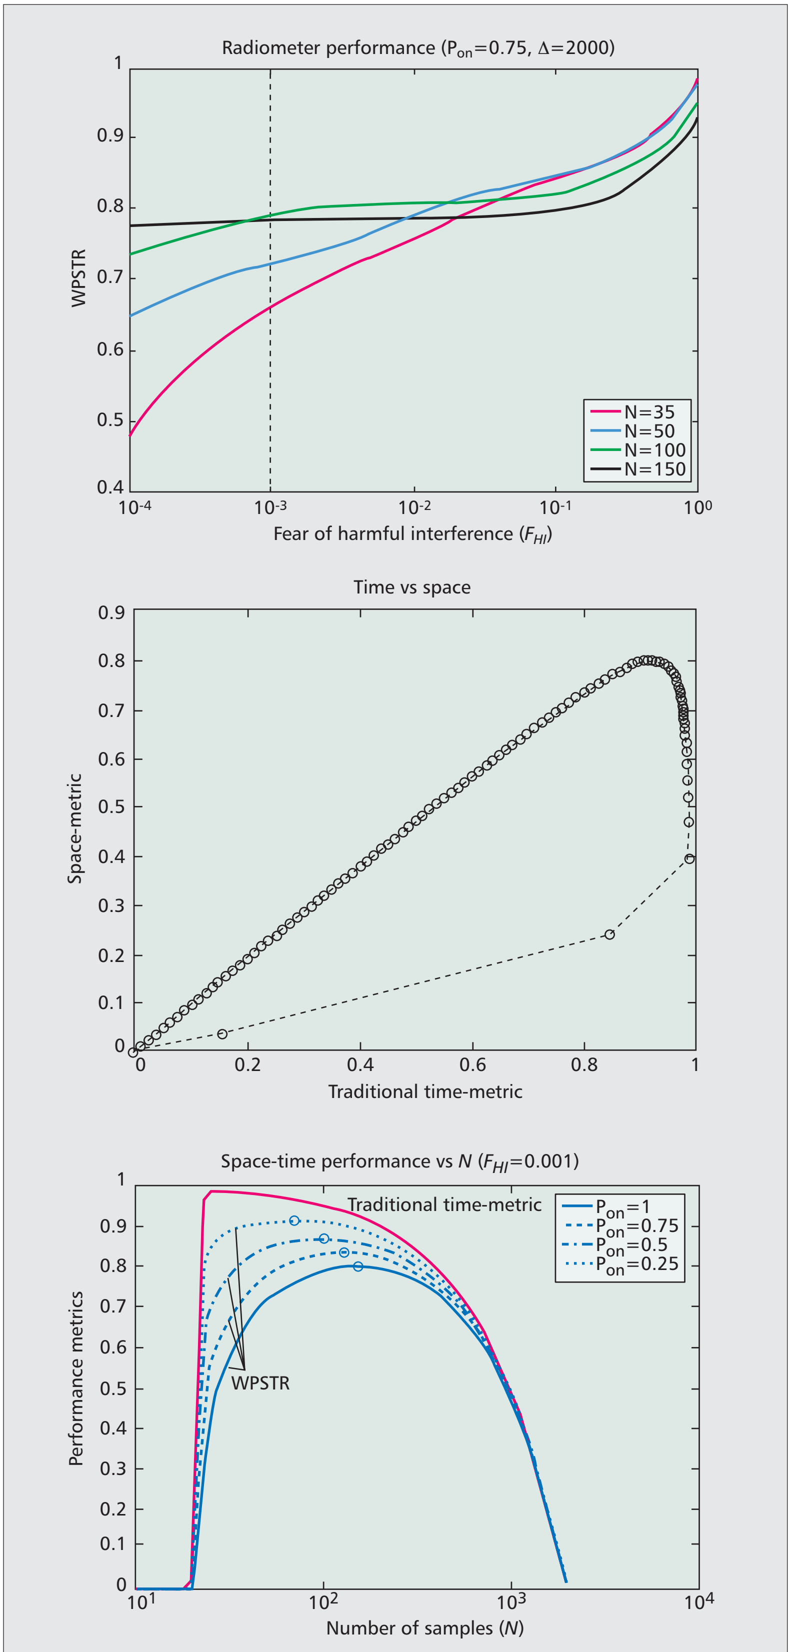

Figure 6. *The plot on top shows the* F_{HI} *vs. WPSTR trade-off for a radiometer as the sensing time* N *varies. Note that the optimal value of the sensing time is a function of the target* F_{HI}*. The plots on the bottom drill down for a particular* F_{HI} *= 0.001 and show the trade-off between the traditional time metric* $\Delta - N/\Delta(1 - P_{FA})$ *and space recovery for a radiometer. The traditional metric underestimates the optimal sensing duration* N *whenever there is a spatial component of the spectrum holes.*

ON/OFF PRIMARIES: SPACE-TIME TRADE-OFF

When the primary signal has both *ON* and *OFF* periods, both space and time must be considered together.

Memoryless Sensing Algorithms — Even this restrictive class of algorithms brings out some interesting trade-offs that are absent in the space-only scenario discussed earlier. The assumptions we make are:

- The detector senses for N contiguous samples and makes a decision based these samples alone. This is clearly suboptimal because longer-term memory could help significantly [11].
- The secondary user's sensing and transmission times are non-adaptive and fixed in advance.
- The primary state does not change within a sensing window. This approximation can lead to a lower rate of missed-detections because if the primary were to turn ON somewhere near the end of the sensing window, there would be a good chance that the detector would not trigger. To counter this, we enforce that the sum of the sensing-duration and the secondary user's transmission-duration is less than the buffer Δ. This ensures that the only way to cause unauthorized harmful interference is by having a sensing error during a window in which the primary is ON.

Numerical Simulations — For simplicity, consider a radiometer facing no fading at all. See [10] for the derivation of expressions for F_{HI} and *WPSTR*, but the results can also be extended to other single-user or even cooperative sensing algorithms. The parameters used to obtain our numeric results are chosen to match those in [12] and are described below:

The TV tower's transmit power is assumed to be $P_t = 10^6$ W, its protection radius $r_p = 134.2$ km, and the no-talk radius $r_n = 150.3$ km. The received power at a distance r from the TV tower is modeled as $P(r) = P_t \cdot l(r)$ where $\log l(r)$ is a piece-wise linear continuous function of $\log r$ chosen to approximate the International Telecommunication Union (ITU) propagation curve given in [12, Fig. 1]. Finally, the exponent in the spatial weighting function $w(r) := A\exp - \kappa(r - r_n)$ is chosen to be $\kappa = 0.02$ km^{-1} .

The top plot in Fig. 6 shows curves depicting the F_{HI} vs. *WPSTR* performance of a radiometer for different sensing times N. It is clear that the optimal N is a function of the desired safety F_{HI} The bottom right plot takes a slice at a fixed F_{HI} = 0.001, and considers the radiometer's *WPSTR* performance as a function of the sensing time. It compares this with the traditional perspective's overall performance metric:

$$\left(1 - \frac{N}{\Delta}\right)(1 - P_{FA}).$$

Notice that at very low N, essentially nothing is recoverable since the F_{HI} forces the detection threshold to be so low that noise alone usually triggers it. There is an optimal value for N that balances the time lost to sensing with the opportunities lost from false alarms, but the traditional perspective is far more aggressive about setting the sensing duration N. This is because the two traditional hypotheses are well separated, but for potential locations close to r_n, the relevant hypotheses are much closer. As illustrated in the bottom left plot in Fig. 6, there is a tension that must be balanced between performance in space (which demands high-fidelity from the radiometer and hence more sample complexity) and the solely time-oriented traditional performance metric.

CONCLUDING REMARKS

It is tempting to force spectrum sensors to be very sensitive so as to guarantee protection to the primary user (e.g., the –114 dBm rule in [5]). But the traditional metrics completely miss that this forces the loss of a significant portion of the spatial spectrum holes because of a presumed lack of diversity. To see the underlying tradeoffs, a new joint space-time formulation is needed that formulates the spectrum-sensing problem as a composite hypothesis test.

Unfortunately, simple single-user strategies cannot obtain enough diversity to get a good trade-off. One needs to look at other sensing strategies like dual detection, collaborative sensing, multiband sensing, and so on, to improve performance. The key is to have a robust way for the secondary user to conclude that it is indeed not deeply shadowed (not being shadowed is, after all, the typical case) and thereby *avoid* being more sensitive than is warranted.

One possibility, that deserves further investigation, is to exploit sensor memory. If a secondary user has seen a strong primary signal in the near past, it knows that it is probably not deeply shadowed. This suggests that cooperative change-detection-based algorithms can improve sensing performance in both space and time.

REFERENCES

[1] M. S. Alouini and A. Goldsmith, "Area Spectral Efficiency of Cellular Mobile Radio Systems," *IEEE Trans. Vehic. Tech.*, vol. 48, no. 4, July 1999, pp. 1047–66.

[2] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Info. Theory*, vol. 46, no. 2, Mar. 2000, pp. 388–404.

[3] S. Huang, X. Liu, and Z. Ding, "On Optimal Sensing and Transmission Strategies for Dynamic Spectrum Access," *IEEE DySPAN*, Oct. 2008, pp. 1–5.

[4] R. Tandra, S. M. Mishra, and A. Sahai, "What is a Spectrum Hole and What Does it Take to Recognize One?" *Proc. IEEE*, May 2009, pp. 824–48.

[5] FCC, "In the Matter of Unlicensed Operation in the TV Broadcast Bands: Second Report and Order and Memorandum Opinion and Order," tech. rep. 08-260, Nov. 2008.

[6] A. Sahai, N. Hoven, and R. Tandra, "Some Fundamental Results on Cognitive Radios," *Allerton Conf. Commun., Control, Comp.*, Oct. 2004.

[7] S. M. Mishra, *Maximizing Available Spectrum for Cognitive Radios*, Ph.D. Dissertation, UC Berkeley, 2010.

[8] F. Baccelli and B. Blaszczyszyn, "Stochastic Geometry and Wireless Networks," *Foundations and Trends in Net.*, vol. 3, no. 3–4, 2009, pp. 249–449.

[9] R. Tandra and A. Sahai, "SNR Walls for Signal Detection," *IEEE J. Sel. Topics Signal Process.*, Feb. 2008, pp. 4–17.

[10] R. Tandra, A. Sahai, and V. Veeravalli, "Space-Time Metrics for Spectrum Sensing," *IEEE DySPAN*, Apr. 2010, pp. 1–12.

[11] A. Parsa, A. A. Gohari, and A. Sahai, "Exploiting Interference Diversity for Event-Based Spectrum Sensing," *IEEE DySPAN*, Oct. 2008.

[12] S. J. Shellhammer *et al.*, "Performance of Power Detector Sensors of DTV Signals in IEEE 802.22 WRANs," *1st Int'l. Wksp. Tech. Policy Accessing Spectrum*, June 2006.

BIOGRAPHIES

RAHUL TANDRA (tandra@gmail.com) is a senior systems engineer at Qualcomm Research Center, Qualcomm Inc., San Diego, California. He currently works on the design and standardization of next-generation WLAN systems. He received his Ph.D. degree from the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley in 2009, where he was a member of the Wireless Foundations research center. In 2006 he worked as a summer intern with the Corporate R&D division of Qualcomm Inc., developing spectrum sensing algorithms for the IEEE 802.22 standard. Prior to that, he received an M.S. degree from Berkeley in 2005 and a B.Tech. degree in electrical engineering from the Indian Institute of Technology Bombay. His research interests are in wireless communication and signal processing. He is particularly interested in fundamental research questions in dynamic spectrum sharing.

ANANT SAHAI [BS'94, SM'96] (sahai@eecs.berkeley.edu) is an associate professor in the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley, where he joined the faculty in 2002. He is a member of the Wireless Foundations center. In 2001 he spent a year at the wireless startup Enuvis developing adaptive software radio algorithms for extremely sensitive GPS receivers. Prior to that, he was a graduate student at the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. His research interests span wireless communication, decentralized control, and information theory. He is particularly interested in spectrum sharing, the nature of information in control systems, and power consumption.

VENUGOPAL VEERAVALLI [S'86, M'92, SM'98, F'06] (vvv@illinois.edu) received his B.Tech. degree in 1985 from the Indian Institute of Technology Bombay (Silver Medal Honors), his M.S. degree in 1987 from Carnegie-Mellon University, and his Ph.D. degree in 1992 from the University of Illinois at Urbana-Champaign, all in electrical engineering. He joined Illinois in 2000, where he is currently a professor in the Department of Electrical and Computer Engineering, a research professor in the Coordinated Science Laboratory, and the director of the Illinois Center for Wireless Systems (ICWS). He served as a program director for communications research at the U.S. National Science Foundation in Arlington, Virginia, from 2003 to 2005. He was with Cornell University before he joined Illinois, and has been on sabbatical at MIT, IISc Bangalore, and Qualcomm, Inc. His research interests include distributed sensor systems and networks, wireless communications, detection and estimation theory, and information theory. He is a Distinguished Lecturer for the IEEE Signal Processing Society for 2010–2011. He has been on the Board of Governors of the IEEE Information Theory Society. He has also served as an Associate Editor for *IEEE Transactions on Information Theory* and *IEEE Transactions on Wireless Communications*. Among the awards he has received for research and teaching are the IEEE Browder J. Thompson Best Paper Award, the National Science Foundation CAREER Award, and the Presidential Early Career Award for Scientists and Engineers (PECASE).

One possibility is to exploit sensor memory. If a secondary user has seen a strong primary signal in the near past, it knows that it is probably not deeply shadowed. This suggests that cooperative change-detection based algorithms can improve sensing performance in both space and time.