Cognitive Radios for Spectrum Sharing: Technical Appendices

Anant Sahai
Mubaraq Mishra
Rahul Tandra
Kristen Ann Woyach

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2008-147
http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-147.html

December 2, 2008

Acknowledgement

We thank the National Science Foundation (grants ANI-326503, CNS-403427, CCF-729122 as well as a Graduate Research Fellowship), C2S2 (Center for Circuit System Solutions), and Sumitomo Electric for their support.

Cognitive Radios for Spectrum Sharing: Technical Appendices

Anant Sahai Shridhar Mubaraq Mishra Rahul Tandra Kristen Ann Woyach

{sahai, smm, tandra, kwoyach}@eecs.berkeley.edu
Dept. of Electrical Engineering and Computer Sciences, U C Berkeley

Abstract

This is a technical appendix to the paper titled *Cognitive Radios for Spectrum Sharing* which will appear in the January issue of the IEEE Signal Processing Magazine. Due the severe space constraints the technical details had to be moved out of the paper and incorporated into this report.

I. Technical details for Figure 1

To calculate the available white space we retrieved the FCC's transmitter database which lists the latitude, longitude, effective radiated power (ERP) and transmitter height above sea level for all licensed transmitters [1]. This database together with the ITU-R recommendations on propagation (ITU-R P-1536-3) were used to calculate the protected radius for all towers. To calculate the value of of the no talk radius, we assumed a 4W cognitive radio with a height above average terrain (HAAT) of 75m. These assumptions are similar to the assumptions made by the IEEE 802.22 working group [2]. We assumed that the total white space was the total number of channels (67 channels – channel 2 to 69 excluding channel 37) times the area of the United States. The available white space was calculated as a percentage of the total white space.

A. Calculating the protected radius

Firstly, we assumed that all signals are ATSC signals i.e. the assumption is made that NTSC signals would switch to ATSC signals with the appropriate transmit power limits as shown in Table I. The required field strength ($E_{r_p}(dBu)$) for the Grade B (protected radius) contour is defined by the FCC for ATSC signals [3] as shown in Table II. For each transmitter we need to determine the distance value at which the received signal is above this specified level for 50% of the locations, 90% of the time. These numbers were calculated using the following procedure which is similar to the procedure used by the IEEE 802.22 Working Group:

- The Effective Radiated Power (dBm) is converted to Effective Isotropic radiated Power (dBW) ($EIRP(dBm)$ = $ERP(dBm)$ + $2.15dB$).
- The Electric field at a distance of 1m from the transmitter is calculated as $E_{1m}(dBu) = 104.8 + EIRP(dBm)$.
- The required path loss is calculated as $RPL = E_{1m} - E_{r_p}$.
- The ITU-R recommendations are used to determine the maximum distance (beyond 1m) at which the path loss is less than (or equal to) RPL for 50% of the locations, 90% of the time. (The ITU-R recommendations provide a mechanism to extrapolate the tables for different distances, heights and frequency [4]).

For a 1MW TV transmitter on channel 38 with a HAAT of 500m the calculated protected radius is 135km.

Channels	*Maximum Transmit Power (ERP) (kW)*
$2-6$	45
$7-13$	160
$14-69$	1000

TABLE I
Maximum transmit power for ATSC signals.

Channels	*Required Field Strength (dBu)*
$2-6$	23
$7-13$	36
$14-69$	$41 - 20\log_{10}(\frac{615}{\frac{f_h+f_l}{2}})$

TABLE II
ATSC (DIGITAL) FIELD STRENGTH DBU REQUIRED AT THE GRADE B CONTOUR. f_l AND f_h ARE THE CHANNEL'S LOWER AND HIGHER FREQUENCY LIMITS (IN MHZ) RESPECTIVELY.

Channels	*Required Field Strength (dBu)*
$2-6$	P(dBm) = E(dBu) -111.8
$7-13$	P(dBm) = E(dBu) -120.8
$14-69$	P(dBm) = E(dBu) $-130.8 + 20\log_{10}(\frac{615}{\frac{f_h+f_l}{2}})$

TABLE III
DBM TO DBU CONVERSION VALUES FOR VARIOUS FREQUENCIES. f_l AND f_h ARE THE CHANNEL'S LOWER AND HIGHER FREQUENCY LIMITS (IN MHZ) RESPECTIVELY. [3]

B. Calculating the no-talk radius

For calculating the no talk radius we need to determine the distance beyond the protected radius where it is safe to transmit. According to [2] the Desired Signal to Interference ratio (SIR) ratio is 12.2dB. Using this value we can calculate the maximum secondary electric field at the protected radius as $E_{r_p} - 12.2$dBu. Using the procedure outlined in Section I-A we can determine the distance beyond the protected radius where a secondary could transmit and have its field strength drop to $E_{r_p} - 12.2$dBu at the protected radius. The only difference is that we use the $F(50, 10)$ curves for predicting the distance. This is to ensure that transmissions from a secondary just outside no-talk radius can cause interference only 10% of the time.

C. Calculating the radius for the -116dBm rule

The -116dBm value is converted to an equivalent dBu value using the equations in Table III. Using the procedure outlined in Section I-A we can calculate equivalent distance. The only difference is that we use the $F(50, 50)$ curves for predicting these distances.

D. Available White Space by population

To calculate available white space in terms of population we used the US Census data of 2000 which lists the population density per zip code [5]. Furthermore, the region occupied by a zip code is specified as a polygon [6]. We assumed that population was uniformly distributed within each zip code. Using this data the total white space was considered to be the population of the continental United States times the number of channels (67 channels).

For every detection rule, the no-talk area around each transmitter is marked out. Using Monte Carlo simulations we can determine the fraction of the area that is available for white space operations. For determining the white space by population, we uniformly sampled the population of the United States. For each sample we then determined the zip code it belonged to and assumed that the sample was uniformly located within the zip code.

II. TECHNICAL DETAILS FOR FIGURE 2

The plot on the top left corner of the figure illustrates the tradeoff between the sensing time N and the operating SNR subject to the constraint that the probability of false-alarm, $P_{FA} = 0.01$, and the probability of missed-detection, $P_{MD} = 0.01$. The dotted curves shows this tradeoff for a radiometer without any noise uncertainty. The green curves shows this tradeoff for a radiometer with noise uncertainty. These curves clearly show that the time-overhead (N) goes to infinity as we approach the SNR wall for the detector. The exact equations describing this tradeoff is given in Equation (5) of [7]. The blue curves show the same tradeoff for a coherent detector.

The plot on the top right corner of the figure illustrates the spatial-overheard for a perfect radiometer ($N = \infty$) in the presence of noise uncertainty. The dotted curve is the tradeoff between $1 - WPAR$ and F_{HI} for a perfect radiometer without noise uncertainty. The exact mathematical expressions describing this tradeoff are given in Equation (27) of [8].

The tradeoff for the case with noise uncertainty is described below. The test-statistic for a perfect radiometer is given by $T^{per} := 10^{\frac{P}{10}} + \sigma_w^2$, where P is the signal power in dB, and σ_w^2 is the actual noise variance, and can lie anywhere between $\frac{1}{\rho}\sigma_{w,nominal}^2$ and $\rho\sigma_{w,nominal}^2$, where $\rho \geq 1$ is the noise uncertainty parameter. Hence, the F_{HI} is computed over the worst-case noise power and is given by

$$F_{HI} = 1 - \mathcal{Q}\left(\frac{\widetilde{\lambda} - \mu(r_n)}{\sigma}\right), \tag{1}$$

where $\widetilde{\lambda} := 10log_{10}(\lambda - \frac{1}{\rho}\sigma_{w,nominal}^2)$, λ is the radiometer detection threshold, σ is the variance of the signal power distribution, and $\mu(r_n)$ is the mean of the signal at the edge of the no-talk radius. Similarly, the probability of finding a hole is given by

$$P_{FH}(r) = 1 - \mathcal{Q}\left(\frac{\widehat{\lambda} - \mu(r_n)}{\sigma}\right), \qquad \forall \quad r > r_n \tag{2}$$

where $\widehat{\lambda} := 10log_{10}(\min(\lambda - \sigma_{w,nominal}^2, 0))$, because for computing $WPAR$, we can assume that we know the noise statistics completely.

For a given value of F_{HI}, we can compute the detector threshold using (1). It is clear that as F_{HI} converges to zero, λ converges to $\frac{1}{\rho}\sigma_{w,nominal}^2$. Hence, for sufficiently low values of F_{HI}, $\lambda < \sigma_{w,nominal}^2$. Tn this case, $\widehat{\lambda}$ in (2) will be $-\infty$, and hence $P_{HI}(r) = 0$ for all $r > r_n$. This forces the $1 - WPAR$ to be one. This behavior can be observed in the blue curve in the plot. There is a minimum threshold for F_{HI}, below which the spatial overhead is 1, that is, no spectrum holes are recovered!

The two plots in the center of the bottom row of the figure show the receiver operating characteristic (ROC) curves for a radiometer at two different operating SNRs (-6 dB and 2.2 dB). The dotted curves in these plots are the ROC curves for the radiometer without any noise uncertainty and for different values of the sensing time N. These curves show that as N increases, both P_{FA} and P_{MD} converge to zero. The solid curves in these plots are the ROC curves for a radiometer with 1 dB of noise uncertainty. For an operating SNR of 2.2 dB, which is above the SNR wall, the ROC curves converge to zero as N increases to infinity. However, for an operating SNR of -6 dB, which is below the SNR walls, the ROC curves converge to 1 as N increases to infinity. This clearly shows: if the operating SNR is below the SNR wall, then increasing N does not improve the error probabilities. Infact, in this regime (SNR below the SNR wall), it is better to use a coin-flipping detector instead of a radiometer.

The equations for P_{FA} and P_{MD} for a radiometer at a given SNR and sensing time N are given by

$$P_{FA} = \mathcal{Q}\left(\frac{\gamma - \rho\sigma_n^2}{\sqrt{\frac{2}{N}}\rho\sigma_n^2}\right) \tag{3}$$

and

$$P_{MD} = 1 - \mathcal{Q}\left(\frac{\gamma - (P + \frac{1}{\rho}\sigma_n^2)}{\sqrt{\frac{2}{N}}(P + \frac{1}{\rho}\sigma_n^2)}\right), \tag{4}$$

where ρ is the noise uncertainty parameter, P is the signal power, σ_n^2 is the nominal noise power, and γ is the radiometer detection threshold. These curves are obtained by first choosing a given P_{FA} and finding a γ that satisfies (3), and then substituting this value of γ in (4), we get the corresponding value of P_{MD}.

The first and the last plot in the bottom row of the figure show the set of CDF's under both hypotheses and an operating SNR of -6 dB and 2.2 dB respectively. In the leftmost plot ($SNR = -6$ dB), the set of CDF's under both hypotheses completely overlap, and hence robust detection is impossible. However, in the rightmost plot ($SNR = 2.2$ dB), the set of CDF's do not overlap, and hence robust detection is feasible in this case. These observations match the observations in the ROC plots described above.

III. Technical details for Figure 3

The leftmost picture in Figure 3 is a cartoon where the values of the mean and standard deviations are representative of numbers we may see in indoor/outdoor scenarios.

The details for obtaining the other three plots can be found in [8]. Here we reproduce some of that discussion.

For M radios with received powers $P_1 \dots P_M$, the Maximum Likelihood (ML) test statistic is the average received power and the detection problem can be stated as:

$$\frac{1}{M}\sum_{i=1}^{M} P_i \underset{No\,Primary}{\overset{Primary}{\gtrless}} \lambda. \tag{5}$$

We can calculate fear of harmful interference (F_{HI}) as follows:
Assuming that $P_i \sim \mathcal{N}(\mu(r), \sigma^2)$, we have $\frac{1}{M}\sum_{i=1}^{M} P_i \sim \mathcal{N}(\mu(r), \frac{\sigma^2}{M})$ where $\mu(r)$ is the mean received power (in dBm) at a distance r from a transmitter with a transmit power of 1MW and a HAAT of 500m. The path loss curve follows the ITU F(50, 90) curve for a center frequency of 615MHz and a height of 500m. The value of σ was 5.5dB.

$$\begin{aligned} F_{HI} &= 1 - \mathcal{P}\left(\frac{1}{M}\sum_{i=1}^{M} P_i \geq \lambda | r_{actual} = r_n\right) \\ &= 1 - \mathcal{Q}\left(\frac{\lambda - \mu(r_n)}{\frac{\sigma}{\sqrt{M}}}\right). \end{aligned} \tag{6}$$

where r_n is the no talk radius. The detector threshold λ must be chosen such that $F_{HI} \leq F_{HI}^{target}$. Hence (6) gives

$$\lambda = \frac{\sigma}{\sqrt{M}}\mathcal{Q}^{-1}(1 - F_{HI}^{target}) + \mu(r_n). \tag{7}$$

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

$$\begin{aligned} P_{FH}(r) &= \mathcal{P}\left(\frac{1}{M}\sum_{i=1}^{M} P_i \leq \lambda | r_{actual} = r\right) \\ &= 1 - \mathcal{Q}\left(\frac{\lambda' - \mu(r)}{\frac{\sigma}{\sqrt{M}}}\right). \end{aligned} \tag{8}$$

The WPAR can be computed as:

$$WPAR = \int_{r_n}^{\infty} P_{FH}(r) w(r)\; r dr, \tag{9}$$

Spatial overhead is specified as $1 - WPAR$. The second plot in Figure 3 shows the spatial overhead versus the fear of harmful interference for various number of cooperating users. As stated in the paper, the -116dBm rule corresponds to a F_{HI} of 2×10^{-4}. For channels 38 the available white space is 55% of the total white space but the -116dBm rule only claims 13.75%. Hence the spacial over head is 0.75. $(1 - \frac{0.1375}{0.55})$. This is the empirical performance of the -116dBm rule.

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

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

The system of cognitive radios causes harmful interference only if every radio individually fails to detect the primary user, and so by independence:

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

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

$$\lambda_{radio,M} = \sigma\mathcal{Q}^{-1}\left(1 - \left[F_{HI}^{target}\right]^{\frac{1}{M}}\right) + \mu(r_n). \tag{10}$$

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

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

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

$$\begin{aligned} P_{FH,system}(r) &= \left(P_{FH,radio}\right)^M \\ &= \left[1 - \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r)}{\sigma}\right)\right]^M. \end{aligned} \tag{11}$$

Substituting (11) in (9), we get the WPAR for the OR rule.

To determine the performance of cooperation with M correlated users, we assumed that $P_1, P_2, \cdots, P_M$ is a jointly Gaussian random vector with marginals given by $P_i \sim \mathcal{N}(\mu(r), \sigma^2)$, where r is the common radial distance from the primary transmitter. Further, the $M \times M$ covariance matrix C has entries $C(i,j)$ given by

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

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

Under this uncertain correlation model it is easy to show that the averaging detector in (5) is the ML detector no matter what the value of ρ is. Further, it is straightforward to show that to meet a low F_{HI} constraint, the averaging detector must design its λ for the worst case correlation, $\rho = \rho_{max}$. For this choice of ρ we have $\frac{1}{M}\sum_{i=1}^{M} P_i \sim \mathcal{N}(\mu(r), \frac{1}{M}[1 + (M-1)\rho_{max}]\sigma^2)$. Therefore,

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

From the above equation we can choose a λ such that the target F_{HI} requirement is met. Given this λ we compute the $WPAR$ performance assuming the nominal model, which corresponds to complete independence, $\rho = 0$, i.e., $\frac{1}{M}\sum_{i=1}^{M} P_i \sim \mathcal{N}(\mu(r), \frac{1}{M}\sigma^2)$. The performance of cooperation for different values of ρ_{max} is shown in the last plot in Figure 3.

IV. Technical details for Figure 4

The technical content of this figure is included in [9]. A short description of that content is included here.

The far left part of the figure is a timeline showing the proposed operation of the our spectrum jail, and the idea of cognitive radio as a band expander. The concept is that spectrum is divided into chunks of utility 1; a cognitive user then has a home band which is worth β of these chunks and can stake this home band against expanding into other bands of utility 1. If it is caught cheating in any of the other bands (using the band when the band's primary is also active), it will be sent to a global jail where it is not allowed to use either the expansion bands or its home band.

The expansion into other bands and into a global jail is depicted in a Markov chain-like structure in the top middle of the figure. Each primary has an independent operating characteristic defined by p and q, the probability of turning off and on, respectively. The cognitive user can respond to this in each band by cheating (with tunable probability P_{cheat}), waiting, seeing a false alarm, or legally transmitting. If caught cheating (with probability P_{catch}), it is sent to a global jail for a sentence of average length $1/(1 - P_{pen})$. It can also be wrongfully convicted of cheating with probability P_{wrong} in each band.

In this game, the cognitive user gains one unit of utility whenever it is cheating or legally transmitting, and gains a utility of β whenever it is not in jail. The cognitive user is trying to use its probability of cheating, P_{cheat} to maximize its average utility—the utility it gets when the Markov chain is in a stationary distribution:

$$\max_{P_{cheat}} U = \max_{P_{cheat}} \pi_{legal} + \pi_{cheat} - \beta\pi_{jail}. \tag{12}$$

where π_{legal} is the stationary probability of legally transmitting and π_{cheat} is the stationary probability of cheating.

From the regulator's perspective, the jail sentence must be set such that the cognitive user has no incentive to cheat. This means that the cognitive user must lose more utility by sitting in jail than it would gain by cheating:

$$BP_{cheat}\pi_{not\text{-}jail} - \beta\pi_{jail} < 0, \tag{13}$$

where π_{not_jail} is the stationary probability of not being in jail, and π_{jail} is the stationary probability of being in jail. This imposes a constraint on P_{pen}:

$$P_{pen} > 1 + P_{wrong} - \frac{\beta}{B}(P_{catch} - P_{wrong}). \tag{14}$$

which goes to 1 as either P_{wrong} or B rise. This makes intuitive sense as when the probability of being wrongfully accused rises, the cognitive user will be sent to jail with the same probability whether cheating or not. It may as well get some utility by cheating. Also, if B is getting larger, the temptation to cheat is rising because the utility the cognitive user is gaining by cheating is greater than it would lose when not able to use its home band. The two top right plots in the figure show how P_{pen} goes to 1 with these parameters for different values of P_{catch}.

From the perspective of the cognitive user, spending time in jail reduces utility. So, if we set P_{pen} to the minimum it needs to be to deter cheating (equality in (14)), the Utility of the Cognitive User plot in the figure shows the steady state utility of secondary user and the fraction of time spent in jail against the expansion B. This plot was created empirically by changing B and calculating the average utility, maximized over P_{cheat} as in (12). In this example, the utility peaks when the cognitive user spends about half of its time in jail.

Again setting P_{pen} to have equality in (14), if we let the cognitive user choose the correct B to maximize its utility, we achieve the plots shown in the bottom right corner.

In the large plot, we assume that P_{wrong} does not scale with the number of bands the cognitive user is expanding into. This is the case in which the cognitive user can be caught in any band (correctly or wrongfully), but there is a particular probability of the cognitive user hearing the "go to jail" command. In this case, the maximal expansion has a closed form:

$$\frac{B}{\beta} = \frac{P_{catch}(1 - P_{TX}) + P_{wrong}(P_{TX} - 2)}{2P_{wrong}(1 - P_{TX})}. \tag{15}$$

The large bottom right plot shows this equation fixing P_{TX} (the probability any one of the primaries is transmitting) and varying P_{catch}. The top cutout shows this same equation but varying P_{catch}.

The second cutout covers the case when P_{wrong} scales with the number of bands the cognitive user is expanding into. This is the case where the cognitive user will hear a "go to jail" command with probability 1, but each primary has their own independent probability of wrongfully sending the cognitive user to jail. This plot was created by numerically optimizing the number of expansion bands when P_{pen} and P_{wrong} are both changing with B.

The bottom middle plot shows the bandwidth expansion and global jail in terms of the overhead such a system would incur. The overhead is defined as

$$Overhead = \frac{\beta + BP_{TX} - (U + \beta)}{\beta + BP_{TX}}, \tag{16}$$

or the fraction of utility the cognitive user is not able to attain because it has to spend time in jail. This plot shows the utility against the bandwidth expansion. Of particular note here is that the maximal expansion never produces an overhead larger than half of the total possible utility, so the cognitive user will never choose to spend more than half its time in jail. Also note that for the maximal bandwidth expansion to be large, P_{wrong} has to be small. In order for the overhead to also be small, P_{wrong} has to be *very* small. This means that this kind of jail system for spectrum will follow the same kind of logic as jail in the real world: it is better to let a thousand guilty men go free than convict one innocent man.

V. Technical details for Figure 5

The calculations for Figure 5 come directly from [10].

The top part of the figure is occupied by an example. Each part of the code illustrated was generated using independent coin flips with a probability $1 - (0.8)^{\frac{1}{3}}$ of being a taboo slot. The resulting AND of the three code parts resulted in an overall overhead of 20%spent on taboo slots.

Below the example, the two plots on the left represent two different slices through the tradeoff between time till conviction, overhead due to taboo slots, and the level of interference that we consider harmful. The formula is (2) in [10] where a central-limit-theorem approximation is used to address this tradeoff for the case of a single secondary system. The background rate of packet-losses for the primary user was set to 5% and the target probability of false conviction was set to 0.005. While this might seem like a very low probability to use the central limit theorem for, in this case the underlying random variables are Bernoulli and so we believe that the CLT is still qualitatively accurate even with the number of time-slots involved in the low thousands. The target probability of missed detection was set at 10%. These parameters were chosen on the basis of the incentive analysis in Figure 4 to correspond to the low overhead required to support a bandwidth expansion of the order of 22: from a single 18MHz home band into the roughly 66 possible television stations that might be in play.

The final plot on the right is from (10) in [10] and represents a lower bound on the time till conviction that has been reinterpreted as a lower bound on the overhead required. There is no closed-form solution for the reinterpretation and so the solutions were computed numerically.

REFERENCES

[1] F. Database, "Summary of TV transmitters." [Online]. Available: http://www.fcc.gov/fcc-bin/tvq?state=&call=&arn=&city=&chan=&cha2=69&serv=&type=3&facid=&list=2&dist=&dlat2=&mlat2=&slat2=&dlon2=&mlon2=&slon2=&size=9

[2] G. Chouinard, "WRAN keep-out region," IEEE 802.22 Meeting Documents, 2006. [Online]. Available: http://www.ieee802.org/22/Meeting_documents/2006_May/22-06-0052-01-0000_WRAN_Keep-out_Region.xls

[3] "Longley-rice methodology for evaluating tv coverage and interference," Federal Communications Commision, OET BULLETIN 69, Feb. 2004.

[4] "Method for point-to-area predictions for terrestrial services in the frequency range 30 mhz to 3 000 mhz," International Telecommunications Commission (ITU), RECOMMENDATION ITU-R P.1546-3, 2007.

[5] U. Census Bureau, "US census 2000 Gazetteer files." [Online]. Available: http://www.census.gov/geo/www/gazetteer/places2k.html

[6] ——, "US Census Cartographic Boundary files." [Online]. Available: http://www.census.gov/geo/www/cob/st2000.html#ascii

[7] R. Tandra and A. Sahai, "SNR walls for signal detection," *IEEE Journal on Selected Topics in Signal Processing*, vol. 2, pp. 4 – 17, Feb. 2008.

[8] R. Tandra, S. M. Mishra, and A. Sahai, "What is a spectrum hole and what does it take to recognize one: extended version," University of California, Berkeley, Berkeley, CA, Tech. Rep., Aug. 2008.

[9] K. A. Woyach, A. Sahai, G. Atia, and V. Saligrama, "Crime and punishment for cognitive radios," *Proceedings of the 46th Annual Allerton Conference on Communication, Control, and Computation*, Sept. 2008.

[10] G. Atia, A. Sahai, and V. Saligrama, "Spectrum Enforcement and Liability Assignment in Cognitive Radio Systems," in *Proceedings of the 3rd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Oct. 2008.