

Codes to Unmask Spectrum Violators in Cognitive Radio Systems

George Atia[†], Venkatesh Saligrama[†] and Anant Sahai[‡]

Abstract—The advent of frequency-agile radios holds the potential for improving the utilization of spectrum by allowing wireless systems to dynamically adapt their spectral footprint based on the local conditions. The gains result from shifting some responsibility for avoiding harmful interference from the static “regulatory layer” to layers that can adapt at runtime. However, this leaves open the major problem of how to enforce/incentivize compliance. We propose a “light-handed” regulation framework where we examine this and focus on assigning liability by detecting the culprits. The key idea is to have the identity of the guilty parties revealed when their signatures appear in the interference pattern. Previously, we only considered a simplistic deterministic scenario [1]. In this paper, we extend the analysis to a more realistic scenario to account for probabilistic collisions, background wireless losses as well as adversarial non deterministic transmission of the guilty parties. We derive converse and achievable bounds on the total system capacity and quantify the fundamental tradeoffs between the various system parameters. The quality of regulatory guarantees is expressed by the time required to convict the guilty, the number of *potential* cognitive systems that can be supported, the number of simultaneously guilty parties that can be resolved and the potential average utilization provided for the secondary systems. We also quantify the tradeoff between stealthiness and actual utilization when culprits purposely trade their utilization to make it harder for the primary system to identify them.

I. INTRODUCTION

The deficiencies in our current model of spectrum allocation originated in the economics/law/public-policy literature in the seminal papers by Coase [2] in 1959 and de Vany, et. al. [3] in 1969 and more recently by Goodman [4]. The focus was mostly on making sure that spectrum was efficiently allocated to the socially most important uses. The work of Mitola [5] introduced the idea of “cognitive radios” that are more intelligent and autonomous than the dumb radios of yesterday. The subsequent Spectrum Policy Task Force report [6] of the FCC generated technical interest in the topic. This report revealed that the inefficiency of the current allocation system went beyond the issue of assigning bands to inefficient uses — the current system has a dramatic underuse of spectrum since much of it is simply not used at all in most places/times. It has been argued that this waste is an inevitable consequence of the current static approach to spectrum access [7].

Dynamic spectrum access thus has the potential for improving the overall utilization of spectrum. Much of the work

has focused on formalizing the concept of spectrum holes and discovering how to utilize them while avoiding causing harmful interference to those primary users that are actually active. The novel idea was to replace static spectral masks as the interface between regulation and implementation and directly deal with limiting interference itself.

Meanwhile, the dominant stream of research has focused on the spectrum sensing aspect of the problem [8]. It has been shown that cooperative sensing could overcome the SNR walls barrier [9] that limit the ability of a single cognitive user to detect severely faded primary users [10], [11]. Although these efforts have led to a new understanding of technical aspects of opportunistic spectrum usage, fundamental gaps remain in translating these efforts to practice. Indeed the very prospect of cooperative spectrum use in turn raises regulatory problems. How should cognitive radios be regulated? When single-user sensing is all that is contemplated, then it is easy to see how the current ultrawideband and unlicensed device paradigm can be extended to cover it: as long as a device senses at an appropriate sensitivity, it is allowed to communicate within an appropriate mask. This is an *a priori* rule that can be relatively easily enforced. But how can we certify the behavior of a network of cooperating users when that network forms dynamically in the field?

What is really missing is a means for *a posteriori* **Spectrum Enforcement** that works in conjunction with some amount of *a priori* certification of the devices themselves. Such a perspective has always been present on the policy side in [2], [3], [4], but has to our knowledge, never been explored technically. A companion paper [12] develops a toy game-theoretic model to discern how effective spectrum enforcement mechanisms must be to properly incentivize cognitive systems to not cheat. Even without such a quantitative model, it is qualitatively obvious that *if there is no chance of being caught, then there is very little incentive to invest serious engineering effort in complying with the regulations*, especially in cases where compliance might result in lower quality of service to the cognitive-radio system’s own users.

A. Motivation

The main idea in this paper is to design a way to trace who is violating a protocol. The analogy is to cars. There are rules designed to protect the safety of others. However, not all of these rules are unbreakable at the device level (e.g. cars do not sense stop signs and force you to stop). Instead, cars are required to have visible license plates that allow violators to be identified and penalized. Our goal is to impose the minimal number of rules that can presumably be checked at device certification time to allow substantial room

[†] G. Atia and V. Saligrama are at the Department of Electrical and Computer Engineering, Boston University, Boston, MA, Emails: {geokamal,sv}@bu.edu

[‡] A. Sahai is at the department of Electrical Engineering and Computer Sciences, University of California Berkeley, CA, Email: sahai@eecs.berkeley.edu

for innovation as new technologies develop. A philosophy of hierarchical punishment was proposed in [1] wherein the offended entity should be able to identify that it is being interfered with and then identify a subset of users who are responsible for causing the interference. There are three potential approaches to ‘identity.’ In the most straight-forward approach, identity is explicitly transmitted by the physical layer as a separate wireless signal in a mandated format. If a primary user experiences harmful interference, then it merely has to decode this signal to learn the identities of all the potential interferers so that they can be penalized. Such an “identification pilot” signal would necessarily impose an overhead on the secondary users and one could analyze the tradeoffs possible. While conceptually simple, this approach has four major shortcomings:

- 1) It forces us to mandate a standard PHY-layer format for transmission of this identity information. This adds additional complexity to systems that want to use another format/modulation for their own signals.
- 2) It imposes a decoder PHY burden on the primary user to implement a way to decode this identity information so that it can penalize secondary systems that are cheating in the vicinity. This is in addition to the primary’s own PHY layer for decoding its own data.
- 3) It does not allow the primary user to distinguish between harmful interference and unfortunate fading or bad luck. A primary user might just be out of range of its transmitter or it might be drowning in harmful interference. There is no way to tell them apart if the secondary identity information was just carried by a separate broadcast signal.
- 4) More subtly, a broadcast identity does not distinguish between the innocent and the guilty. Thus it greatly reduces the incentive to deploy innovative approaches to reduce interference. For example, a cognitive-radio network might be able to use beamforming to null out its transmissions at the primary receiver. However, if any other cognitive radio causes harmful interference, the careful radios in the neighborhood will also be punished.

The second approach to identity trades beacon overhead for reporting overhead. Cognitive radios could be required to keep detailed records of their trajectories and operations. These logs could then be regularly uploaded to the authorities and searched to find the culprits whenever a credible report of harmful interference is filed. This avoids the first two shortcomings, but does nothing about the second two.

Because of these issues, we explore a different approach in which the identity of a device is implicitly announced by the pattern of use/interference itself. For simplicity we assume that all wireless nodes have access to a common sense of time and can divide both time and frequency into slots of moderate length. The complete system consists of a set of frequency-specific binary codes on users that govern their medium access to that frequency band. This restriction is hard-coded into the device and verified as such during the device certification process. In addition, there might be a feedback path by which nodes can be told to “cease and desist” their interfering activities, but this is not studied here.

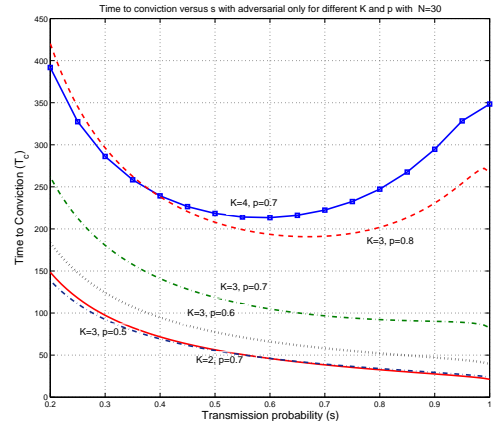


Fig. 1. Time to conviction T_c versus violator’s transmission probability s with $N = 30$, $K = 2, 3, 4$ and utilization 50 – 70%

The focus here is entirely on being able to identify the culprits. Our paper [1] introduced this problem and showed how codes can be arranged so that the primary can rule out “natural causes” (e.g. shadowing) for degraded performance and decide that foul play must be at work. The paper also dealt with liability assignment using the idea of superimposed codes providing bounds on fundamental tradeoffs between the various system parameters within an even more idealized context wherein there was effectively no noise. In this paper, we extend our analysis to a more realistic setup where guilty parties can use probabilistic transmission to make it harder for the primary system to reveal their identities. We also account for background wireless losses as well as non deterministic collision models. In Section II we provide the problem setup. Section III provides converse bounds on the total system capacity. Achievable rates are presented in Section IV and Conclusions are provided in V.

II. PROBLEM SETUP

A primary system observes a degradation in performance, i.e. packet losses, due to spectrum violation. The goal is to identify the guilty parties. Every cognitive user is assigned a different binary codeword that defines its allowable transmission slots and this code is known to the enforcement system. Users are only allowed to transmit during the slots where their codes are equal to 1. Even though a TDMA strategy would be simple and would guarantee enforcement, it is significantly inefficient in terms of utilization as argued in [1]. One important quantity of interest is how long does it take to identify the guilty parties. We define this as the time to conviction T_c . The second important aspect is how powerful the code is in terms of how many cognitive systems it supports and how large the size of the guilty set could be. The goal is to identify the guilty parties by matching the interference pattern to the known collection of codewords. In the next sections we quantify fundamental tradeoffs between the time to conviction T_c , the total number of potential users N , the size of the culprit set K and the average secondary utilization p . The i -th cognitive system is assigned a binary codeword c_i of length T_c . The collection of such codewords

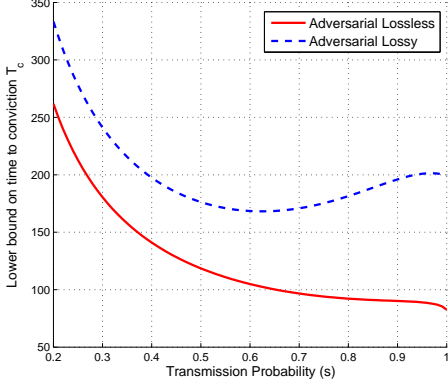


Fig. 2. Time to conviction T_c versus violator's transmission probability s for lossy and lossless models. $N = 30$, $K = 3$, $\alpha = \beta = 10\%$, Utilization=70%.

is denoted C . Let g denote the set of guilty users and y denote the observed binary vector of packet drops. In [1] we considered a best case scenario with zero background noise level. Primary packets were always dropped whenever there is an illegal transmission and received with probability 1 during periods of no interference. Furthermore, in [1] a culprit transmits whenever he is allowed to transmit, hence the behavior of the guilty parties is deterministic as their transmission times coincide with the instants when the code of *any* of its users is equal to 1. The key difference in this paper, is considering a more realistic scenario where we account for random collisions, background losses as well as non deterministic transmission of the guilty users. We quantify fundamental tradeoffs between the various system parameters, namely, the time to conviction T_c , the total number of systems N , the size of the guilty set K and the potential secondary utilization p for the probabilistic loss and interference models. We derive both converse and achievable bounds on the total system capacity of the detect-violator-coalition problem.

III. CONVERSE

In this section we prove a lower bound on the time to conviction T_c . The channel is assumed to be memoryless. In [1], as we considered a deterministic interference/transmission model, the signature of the culprits was perfectly mapped to the interference pattern. Now we would like to account for probabilistic collisions, background wireless losses and also non deterministic transmission patterns of the culprit users. By probabilistic collisions we refer to the fact that, with non zero probability, primary packets could still be successfully received even with an illegal concurrent transmission. Background losses are packet drops due to randomness in the wireless environment such as fading and shadowing.

A. Adversarial probabilistic transmission

Next we consider the case where violators transmit with a known probability s whenever transmission is possible. By doing so, the culprits choose to tradeoff their utilization

hoping to make the identification task harder. Their transmission is assumed to be independent across open slots. For now the only source of randomness is due to this adversarial transmission protocol. In the next subsection we also consider misses, due to non deterministic collisions, and false alarms due to background noise and wireless losses. Note that if violators use s which is too low, their effective utilization will be substantially reduced, hence there is an incentive not to make it too small to tradeoff stealthiness with utilization. The following analysis assumes a uniform distribution for the culprits and that at each instant of time there are ℓ potential users (i.e. ℓ ones per row of C). The entropy of the random vector g can be written as:

$$\begin{aligned} H(\mathbf{g}) &= \log \binom{N}{K} \stackrel{(a)}{=} H(\mathbf{g}|C, s) - H(\mathbf{g}|\mathbf{y}, C, s) \\ &= I(\mathbf{g}; \mathbf{y}|C, s) = H(\mathbf{y}|C, s) - H(\mathbf{y}|\mathbf{g}, C, s) \end{aligned} \quad (1)$$

(a) follows from the independence of g and C and from the fact that $H(\mathbf{g}|\mathbf{y}, C, s) = 0$ to be able to identify the guilty set g given the observation y and the codes C . Now we compute each of the terms above :

$$\begin{aligned} H(\mathbf{y}|C, s) &\leq T_c H(\mathbf{y}_1|c_1, s) \\ &= T_c [-p(y_1 = 1|c_1, s) \log p(y_1 = 1|c_1, s) \\ &\quad - p(y_1 = 0|c_1, s) \log p(y_1 = 0|c_1, s)] \end{aligned} \quad (2)$$

Note that: $y_i = u(c_i \tilde{g})$, with $\tilde{g} = g \wedge a_i$, where a_i is a binary activity vector representing whether the culprit is active or not. In other words, $a_i(n) = 1$ with probability s , if $(n \in G$ and $c_i(n) = 1)$, and 0 otherwise. Thus for $N - \ell \geq K$,

$$\begin{aligned} p(y_1 = 1|c_1, s) &= \Pr[u(c_1 g) = 1] \cdot \Pr[y_1 = 1|u(c_1 g) = 1, s] \\ &\stackrel{(c)}{=} \Pr[u(c_1 g) = 1] \\ &\times \sum_j \Pr[u(c_i \tilde{g}) = 1|u(c_1 g) = 1, j] \Pr[j|u(c_1 g) = 1] \\ &= \left[1 - \frac{\binom{N-\ell}}{\binom{N}{K}} \right] \sum_{j=1}^{\min\{\ell, K\}} (1 - (1-s)^j) \frac{\binom{\ell}{j} \binom{N-\ell}{K-j}}{\binom{N}{K} - \binom{N-\ell}{K}} \end{aligned} \quad (3)$$

In (c) we condition on the number of culprits j whose codes allow them to actually transmit. Denoting the binary entropy by $H_2(\cdot)$:

$$\begin{aligned} H(\mathbf{y}|C, s) &\leq T_c H_2 \left(\left[1 - \frac{\binom{N-\ell}}{\binom{N}{K}} \right] \sum_{j=1}^{\min\{\ell, K\}} (1 - (1-s)^j) \frac{\binom{\ell}{j} \binom{N-\ell}{K-j}}{\binom{N}{K} - \binom{N-\ell}{K}} \right) \\ &= H_2(\eta_1) \end{aligned} \quad (4)$$

Now we compute the second entropy term in Eq.(1):

$$\begin{aligned} H(\mathbf{y}|\mathbf{g}, C, s) &\stackrel{(d)}{=} \sum_{i=1}^{T_c} H(y_i|\mathbf{g}, c_i, s) \\ &= \frac{1}{\binom{N}{K}} \sum_{i=1}^{T_c} \sum_g H(y_i|g, c_i, s) \\ &= \frac{1}{\binom{N}{K}} \sum_{i=1}^{T_c} \sum_{j=0}^{\min\{\ell, K\}} \binom{\ell}{j} \binom{N-\ell}{K-j} H_2((1-s)^j) = T_c \gamma_1 \end{aligned} \quad (5)$$

Equality (d) follows from the fact that $y_i = u(c_i \tilde{g})$ which, conditioned on g , is independent of $y_k \forall k \neq i$. Note that the internal sum in the last summation does not depend on i . This is due to the uniform distribution assumption of g and the fixed number of ones per row of C . Thus a lower bound on T_c can be written as:

$$T_c \geq \frac{\log \binom{N}{K}}{H_2(\eta_1) - \gamma_1} \quad (6)$$

Figure 1 shows a lower bound on the time to conviction T_c as a function of the transmission probability s for different values of K and utilization levels p . These bounds suggest that for certain values of K and p , if the guilty users reduce their transmission probability that could lead to a longer time till identification at the expense of a reduced effective utilization. As we show later, this effect is exhibited in the derived achievable bounds in Section IV.

B. Probabilistic transmission and stochastic losses model

Here we also include miss and false alarm probabilities due to probabilistic collisions and wireless losses, respectively. Let α and β denote the (conditional) false alarm and miss probabilities, respectively. We assume that misses and false alarms are independent across time slots and misses are independent of the number of active users. Modifying the equations above it is not hard to show that:

$$\begin{aligned} & \Pr[y_1 = 1 | c_1, s, \alpha, \beta] \\ &= (1 - \beta) \Pr[u(c \cdot \tilde{g}) = 1] + \alpha \Pr[u(c \cdot \tilde{g}) = 0] \\ &= \alpha + (1 - \beta - \alpha) \left[1 - \frac{\binom{N-\ell}{K}}{\binom{N}{K}} \right] \sum_{j=1}^{\min\{\ell, K\}} (1 - (1-s)^j) \frac{\binom{\ell}{j} \binom{N-\ell}{K-j}}{\binom{N}{K} - \binom{N-\ell}{K}} \end{aligned} \quad (7)$$

Thus,

$$\begin{aligned} H(\mathbf{y}|C, s) &\leq T_c H_2 \left(\alpha + (1 - \beta - \alpha) \left[1 - \frac{\binom{N-\ell}{K}}{\binom{N}{K}} \right] \right) \\ &\times \sum_{j=1}^{\min\{\ell, K\}} (1 - (1-s)^j) \frac{\binom{\ell}{j} \binom{N-\ell}{K-j}}{\binom{N}{K} - \binom{N-\ell}{K}} = T_c H_2(\eta_2) \end{aligned} \quad (8)$$

Similarly,

$$\begin{aligned} H(\mathbf{y}|\mathbf{g}, C, s) &= \frac{1}{\binom{N}{K}} \sum_{i=1}^{T_c} \sum_{j=0}^{\min\{\ell, K\}} \binom{\ell}{j} \binom{N-\ell}{K-j} \\ &\times H_2((1-\alpha)(1-s)^j + \beta(1-(1-s)^j)) = T_c \gamma_2 \end{aligned} \quad (9)$$

A lower bound on the time to conviction can now be readily written as:

$$T_c \geq \frac{\log \binom{N}{K}}{H_2(\eta_2) - \gamma_2} \quad (10)$$

Fig. 2 shows lower bounds on the time to conviction for both the lossy and the lossless models.




Fig. 3. Time to Conviction versus Effective utilization. Guilty users make it harder to get caught by reducing their transmission probability at the expense of a reduced effective utilization. In this setup $N=100$, Original utilization $p = 65\%$.

IV. ACHIEVABLE BOUND FOR CAPACITY

Next we prove an achievable rate using random coding. Codewords of length T_c are independently and randomly generated using bernoulli p distribution. A coalition of K codewords corresponds to a K -user culprit set. Given the output of the channel y we select a set of K users and declare them as the guilty set. Denote a coalition of codewords of size w by $X_{(w)}^{T_c}$ which is T_c independent copies of a collection of vectors $X_{(w)}$. The coalition of K codewords and the observed packet drops y^{T_c} have a joint distribution $p(X^{T_c}, Y^{T_c}) = \prod_{t=1}^{T_c} p(x_{(w)}(t)p(y(t)/x_{(w)}(t))$. Decoding is achieved using typical set decoding. An error occurs if the generating set of codewords and the channel output y are not typical or if any set other than the guilty set and y are jointly typical. We define the error event E_i as follows.

E_i : The event that a set which differs in exactly i users from the true culprits set is jointly typical with y . The probability of such an event is denoted $P(E_i)$. Using the union bound the average probability of error can now be bounded as:

$$P_e \leq P(E_0^c) + \sum_i P(E_i) \quad (11)$$

The first term corresponds to the event that the output y and the input to the MAC channel (i.e. the transmitted codewords of the true coalition) are not typical. It is not hard to show that the probability of this event goes to zero for sufficiently large T_c [13]. Before we prove the achievable rate we prove the following lemma:

Lemma 4.1: If $X_{(K)}$ and Y are jointly distributed as $p(x_{(K)}, y)$ then the probability that a set of K codewords $\tilde{X}_{(K)}^{T_c} = (\tilde{X}_{(i)}^{T_c}, X_{(K-i)}^{T_c})$ formed by replacing i of the codewords of the true coalition, with i independent codewords (hence also independent of Y^{T_c}) and of the remaining $K-i$ codewords, with the same marginals, is jointly typical can be bounded by:

$$\Pr[(\tilde{X}_{(i)}^{T_c}, X_{(K-i)}^{T_c}, Y^T) \in A_\epsilon^{T_c}] \leq 2^{-T_c I(X_{(i)}; X_{(K-i)}, Y)} \quad (12)$$

where $A_\epsilon^{T_c}$ is the set of jointly typical sequences $(X_{(K)}^{T_c}, Y^{T_c})$ [13].

Proof:

$$\begin{aligned}
\Pr[(\tilde{X}_{(i)}^{T_c}, X_{(K-i)}, Y^T) \in A_\epsilon^{T_c}] &= \sum_{x_{(K)}^{T_c}, y^{(T_c)} \in A_\epsilon^{T_c}} P(x_{(i)}^{T_c}) P(x_{(K-i)}^{T_c}, y^{T_c}) \\
&\leq 2^{T_c H(X_{(K)}, Y)} 2^{-T_c H(X_{(i)})} \cdot 2^{-T_c H(X_{(K-i)}, Y)} \\
&= 2^{-T_c I(X_{(i)}; X_{(K-i)}, Y)} \quad (13)
\end{aligned}$$

Intuitively this means that the probability that replacing i codewords of the coalition with i independent codewords would result in jointly typical sequences scales exponentially with the of negative the mutual information between i codewords of the coalition and the remaining $K - i$ codewords and the output. Now we can bound $P(E_i)$ by summing over all sets that differ from the true coalition in i codewords. We can form as many as $\binom{K}{i} \binom{N-K}{i}$ such sets. Thus,

$$P(E_i) \leq \binom{K}{i} \binom{N-K}{i} 2^{-T_c I(X_{(i)}; X_{(K-i)}, Y)} \quad (14)$$

We can now bound the total error probability for T_c sufficiently large as :

$$\begin{aligned}
P_e &\leq \sum_i P(E_i) \\
&\leq \sum_{i=1}^K \binom{K}{i} \binom{N-K}{i} 2^{-T_c I(X_{(i)}; X_{(K-i)}, Y)} \\
&\leq N^K 2^{-T_c \min_i \{I(X_{(i)}; X_{(K-i)}, Y)\}} \quad (15)
\end{aligned}$$

Hence the rate $\frac{K \log N}{T_c} \leq \min_i \{I(X^{(i)}; X^{(K-i)}, Y)\}$ is achievable. Now using the data processing inequality it's not hard to show that: $\frac{K \log N}{T_c} \leq \min_i \{I(f(X^{(i)}); f(X^{(K-i)}, Y))\}$. We can choose the function $f(\cdot)$ to be the integer sum of these sets, i.e. the number of users with potential transmission in a given slot among a given group of users.

Fig. 3 shows how probabilistic transmission could make it harder to get caught, i.e. longer time to conviction is needed, at the expense of a reduced effective utilization since on average culprits end up using $s \cdot p$ of their potential resources. In Fig. 4 we show upper and lower bounds on time to Conviction versus total number of users N with probabilistic losses (due to non deterministic collisions and background noise) and probabilistic transmission.

V. CONCLUSIONS

The main results of this paper is the quantification of the fundamental tradeoffs between the various system parameters under realistic transmission and loss models. We provide both upper and lower bound on the total system capacity expressed in terms of how many coalitions of size K out of N total number of users could be identified within time to conviction T_c . It is shown that: 1) Supporting a larger number of potential users and increasing robustness come at the expense of increased time till conviction; 2) A fundamental tradeoff exists between efficiency, in terms of achievable utilization rates, and timeliness; 3) Allowing a very large number of distinct identities for potential cognitive users




Fig. 4. Upper and lower bounds on time to Conviction versus total number of users N with losses and probabilistic transmission. In this setup, by reducing s secondary users make it harder for the primary to identify them (longer T_c) at the expense of a lower effective utilization. In this figure Utilization $p = 70\%$, $K = 3$.

becomes prohibitively expensive suggesting the need for a gradual punishment mechanism wherein innocent bystander systems might incur short periods of false conviction; 4) Under certain scenarios, revealing the spectrum violators' identities might come at a slightly higher cost (longer T_c) if the guilty parties tradeoff their utilization by adopting randomized probabilistic transmission.

REFERENCES

- [1] George Atia, Anant Sahai, and Venkatesh Saligrama, "Spectrum enforcement and liability assignment in cognitive radio systems," in *Proceedings of the IEEE DySPAN*, Chicago, IL, Oct. 2008.
- [2] R. H. Coase, "The federal communications commission," *The Journal of Law and Economics*, vol. 2, pp. 1–40, Oct. 1959.
- [3] A. De Vany, R. D. Eckert, C. T. Meyers, D.J. O'Hara, and R. C. Scott, "A property system for market allocation of the electromagnetic spectrum: A legal-economic-engineering study," *Stanford Law Review*, vol. 3, pp. 145–162, 1969.
- [4] E. Goodman, "Spectrum rights in the telecom to come," *San Diego Law Review*, vol. 41, pp. 269–404, 2004.
- [5] Joseph Mitola, *Cognitive Radio: An integrated agent architecture for software defined radio*, Ph.d. thesis, KTH Royal Inst. of Tech., Stockholm, Sweden, 2000.
- [6] FCC, "FCC spectrum policy task force report 04-113," [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/, May 2004.
- [7] Rahul Tandra, Shridhar Mubaraq Mishra, and Anant Sahai, "What is a spectrum hole and what does it take to recognize one?," *Accepted to the Proceedings of the IEEE, special issue on Cognitive Radio*, 2008.
- [8] Anant Sahai, Niels Hoven, and Rahul Tandra, "Some fundamental limits on cognitive radio," in *Forty-second Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2004, IEEE.
- [9] Rahul Tandra and Anant Sahai, "SNR walls for signal detection," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 2, no. 1, pp. 4–17, Feb. 2008.
- [10] Shridhar M. Mishra, Anant Sahai, and Robert W. Brodersen, "Co-operative sensing among cognitive radios," in *Proceedings of the International Conference on Communications (ICC)*. IEEE, June 2006, vol. 4, pp. 1658–1663.
- [11] George Atia, Erhan Ermis, Shuchin Aeron, and Venkatesh Saligrama, "Robust energy efficient cooperative spectrum sensing in cognitive radios," in *Proceedings of the forty-fifth Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sept. 2007.
- [12] Kristen A. Woyach, Anant Sahai, George Atia, and Venkatesh Saligrama, "Crime and punishment for cognitive radios," in *Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sept. 2008.
- [13] Thomas M. Cover and Joy A. Thomas, *Elements of Information Theory*, New York: John Wiley and Sons, Inc., 1991.