

# Can we incentivize sensing in a light-handed way?

Kristen Woyach, Padmini Pyapali, and Anant Sahai  
Department of EECS

University of California at Berkeley

Email: kwoyach@eecs.berkeley.edu, padmini.pyapali@gmail.com, sahai@eecs.berkeley.edu

**Abstract**—The current approach to regulate spectrum sensing seems light-handed (just specify the target sensitivity), but has very heavy-handed consequences. The problem is that sensitivity is an intermediate metric that is convenient for certification, but it does not reflect the true externalities imposed by inadequate sensing. This mismatch causes overhead: usually an overinvestment in sensing at the single-radio level and the inability to exploit synergies across multiple radios because the benefit does not show up in the certification metric of sensitivity. Light-handed approaches have the benefit of allowing such flexibility but come with their own overheads. This paper explores this question in some detail to see if criminal-law inspired light-handed approaches can work without imposing more overhead than they are worth.

## I. INTRODUCTION

How would one build a system that can incentivize radios to follow spectrum sharing rules? What does it even mean to “incentivize” man-made devices which are engineered to interact in particular ways? At first glance, using light-handed, policing-based regulation to enforce good behavior in engineered systems may seem unnatural, but in the context of spectrum sharing and dynamic spectrum access, it may be precisely what is required.

It is largely agreed that the current command-and-control system for regulating spectrum access is wasteful [1], evidenced by the fact that only approximately 5% of the available spectrum is actually being used at any given location at any given time [2]. Although the solutions to this inefficiency vary widely, the proposals share a common theme: flexibility. The core problem is that the current system cannot appropriately match available bandwidth to application needs. So, most spectrum sits allocated but idle. Meanwhile, new entrants who may only want small amounts of bandwidth for sporadic periods of time must wait for large swaths to become available and then pay billions for an allocation they cannot completely fill at all times and all places.

Despite its inefficiency, the current system is particularly good at the sister problem to allocation: enforcement. Having an exclusive license is useless if you have no guarantee against harmful interference from unlicensed systems. In the current model, the threat of harmful interference is moot: if chunks of spectrum are allocated to specific companies, and tower characteristics are fixed and known, enforcement is simple. If a license holder experiences interference, the FCC can simply pull out a directional antenna to determine where the interference is originating, assign liability to the party owning the offending system, and subsequent legal action will take care of the problem.

As devices become more agile, trading flexibility for simple enforcement is no longer as appealing. Unfortunately, even conceptualizing enforcement for agile devices is difficult. Consider, for example, the “hit and run radios” introduced by Faulhaber [3]: if radios can dynamically choose their frequencies and positions, they can presumably turn on in one space/frequency location, cause considerable harm, and then turn off and move. Such operation, even if inadvertently harmful, would grant the user the radio equivalent of Plato’s Ring of Gyges: the radio could never be tracked, could never be punished, and so could do whatever it pleased without fear of reproach.

What kind of approach, then, will allow a better tradeoff between flexibility and enforcement? The first step is to understand the tools at our disposal. We must understand that the solution to any enforcement problem could exist at several different time-scales: regulation/certification time, human runtime and device runtime.

Whether the new form of regulation looks like a real-time market, a commons, or a system of dynamic overlays and underlays, devices will always need to be certified by the FCC before they can be deployed. Therefore, there exists the regulation/certification time-scale at which we define the basic operating requirements for any deployed device. Rules must be relatively static, so the regulation time scale is much longer than any other time-scale of interest. Many of the current engineering approaches to enforcement consist of rules that are defined and policed exclusively at certification time. For example, the FCC currently requires that devices coexisting through sensing in the TV bands must be able to demonstrate at certification time that they can sense a primary at a sensitivity of -114dBm [4]. As another example, devices with policy decision engines [5] are intended to prove compliance only at certification time. No further policing is assumed in either case.

Unfortunately, restricting enforcement to certification alone will likely incur a prohibitive amount of overhead. It is shown in [6]–[8] that a single radio sensing alone must give up a significant fraction of its potential transmit opportunity (in either time or space) looking for a primary. Cooperative schemes with multiple radios participating in the sensing process can lower this overhead. Unfortunately, a network of radios working together to sense is difficult, if not impossible, to certify because the relevant operating conditions cannot all be tested at certification time [9]. We offer an alternative approach in this paper that allows simple certification by moving some of the enforcement requirements closer to the

time of use.

Once a system is certified, any remaining disputes must be handled at runtime. But runtime, too, can be split into two classes: human runtime and machine runtime. Human runtime allows disputes to be resolved at the level of Coasian bargains [10] or lawsuits [11]. It is presumed in much of the spectrum policy literature that any disputes not specifically resolved with rules can be resolved by courts (see, for example, the argument in [12]).

However, human runtime is likely insufficient on its own. Court decisions, or Coasian bargains, are very good for solving disputes between companies or causing large-scale shifts in manufacturing. But what about small disputes? Flexible, mobile devices may make mistakes or cause harm on the level of individual devices, not on the level of whole service providers. Suing a company over a few minutes of harmful interference is impractical. If we want to be able to account for small infractions, some policing must be done at the machine runtime. Unfortunately, policing at this time-scale is difficult [13]. It requires, at the very minimum, a method of determining identity and a punishment system. We will assume that a method for determining identity, like the one described in [14], is included in the certification-time requirements so we can focus on the policing and deterrence aspects of the problem.

It is tempting to wonder whether radios could self-enforce at runtime.<sup>1</sup> Is the law of the jungle enough to bring a productive order to wireless? In [16], it is shown that if radios could cause each other a similar amount of interference, the radios could indeed self-enforce and come to one of a range of stable and fair equilibria. However, real systems operate at different powers, have different communication ranges, and have different sensitivities to interference. Real radios would not necessarily obtain a socially-advantageous equilibrium. This problem is coined by Faulhaber as the “power-mix problem” [17] and exists in other literature as a problem of heterogeneity [18]. We must engineer and impose sharing rules beyond the law of the jungle in order to allow heterogeneous systems to coexist. Since direct certification of compliance is likely costly, the goal here is to understand how to achieve desirable behavior with a light-handed policing scheme.

But how do we even think about incentive-based enforcement for insentient devices that have no inherent preferences? In the human realm, considerable work has been done in understanding the economics of law. According to [19], the original conceptions of an economic theory of law were developed by Montesquieu [20], Beccaria [21], and Bentham [22], and later revitalized by Becker [23]. The field of economics of law and punishment has developed significantly since Becker’s work; the general theory can be found in [19], [24]–[26]. We propose that the ideas from this body of work can be mapped into the machine realm to develop a kind of “cyber-justice” in which manufacturers are encouraged to build devices that follow sharing rules by imposing appropriate sanctions *on the*

<sup>1</sup>This is the basic argument employed by commons proponents [15]: set maximum power limits for devices and they will not cause much interference. The interference they do experience will eventually be mitigated by better device designs.

*radios themselves* for non-compliance.

The traditional viewpoint for understanding the effect of punishment is utilitarian: a person will commit a crime if they can expect to receive a positive utility.<sup>2</sup> Philosophically, this view works for humans – we have built-in experiences of pleasure or pain that can be mathematically abstracted by utility functions [22].

Radios, on the other hand, feel neither pleasure nor pain. Their utility therefore must be derived from the desires of their human masters. Unfortunately, it is difficult to even unambiguously specify a master. Is it the current owner or the original designer?<sup>3</sup> Take advertising, for example: to the customer, the content they want is the product, the content provider is the supplier, and advertising is an annoyance. But from the content provider’s perspective, the users may be the product, the content is the bait, and the advertiser is the customer! Who is the master of the device and what is the utility function?<sup>4</sup>

Perhaps we should define the radio’s master as the entity that is liable for the radio’s actions. If the sanction comes in the form of a fine, the master is the one responsible for paying it. But who should pay for a radio’s misconduct? The user who tried to do something illegal or the manufacturer who gave the user the power to do something illegal? Even liability is difficult to define.

Perhaps, then, a different approach is warranted. In the human realm, every person may value things differently, but money to support a life-style and the freedom to live it are somehow universal. Therefore, fines and imprisonment are natural and universal punishments applicable to every human being.<sup>5</sup>

Ideally, we would like to find similarly universal sanctions for radios. As a first attempt, consider applying the human sanction of fines to radios. Radios do not naturally have a need for money, so applying a monetary fine would require significant external infrastructure which may or may not be worth the cost to implement. Consider the two competing visions of the future of technology: convergence or ubiquitous computing [30]. In the former, all functionality you need is available from a single wireless device (the extreme case of Internet, camera, GPS, and voice all available on a cell phone) In this case a billing system may already be present in the device and the outside infrastructure required for a monetary sanction might require little extra cost.

<sup>2</sup>In order for this utilitarian viewpoint to hold, we must assume risk-neutrality, a positive utility for the crime itself, a probability of getting caught, and some kind of sanction.

<sup>3</sup>The radio’s utility function likely should be thought of as an uncertain function in two dimensions: the dimension representing the user’s desires, and the dimension reflecting the designer’s desires. In our companion paper on spectrum zoning, [27], we consider decisions based on uncertain utility functions as a robust optimization problem. Creating truly universal deterrents based on this uncertain utility function is a very important open problem.

<sup>4</sup>This discussion even assumes that the device serves a self-interested master. A radio whose purpose is anti-competitive may have a utility function based on its ability to disrupt its neighbors. For a particularly amusing example of the complications arising from preferences based on others’ valuations, see [28]. In this paper, we will consider greedy, but not malicious radios.

<sup>5</sup>These universal currencies could also be thought of as “primary goods” in the language of Rawls [29].

In the ubiquitous computing future, computing resources are distributed among appliances, piece of furniture, etc, forming a large network of resources. This network is unlikely to have central billing, so creating a new billing system for the sole purpose of fining non-compliant radios would likely be expensive. It would be particularly unnecessary if the radio itself has a more easily accessible, inherent currency.

By definition, a radio is built to transmit and receive signals; it has a predominantly communicative purpose. By virtue of it needing to operate in the real world, the radio is subject to the laws of physics and information theory which determine the limits on how quickly the radio can transmit information and how much power and how many degrees of freedom<sup>6</sup> are required to do so. The radio therefore has limitations on its ability to communicate through limitations on energy (the capacity of its battery, for example), and how many degrees of freedom it can occupy. These two factors can be thought of as a radio's universal currencies, just like money and freedom are universal currencies for humans.

In this paper, we use the degree-of-freedom currency to investigate a theory of economics for radio law by looking at how well different sanctions can deter radios from cheating on sharing rules. For this investigation, we define a sharing rule to govern the interaction between a priority user, or primary, and a non-priority, or cognitive, user: the cognitive user must not transmit when the primary is also transmitting.<sup>7</sup> We then use a light-handed, punitive scheme to encourage a cognitive user to sense for the presence of a primary and avoid transmitting at the same time. Sensing is assumed to be expensive (in a way that we will define precisely later) to the cognitive user, so the deterrence system must get the cognitive user to choose to sense despite the cost of doing so.

We begin by modeling the problem in Section II. Section III uses fines to establish a baseline for discussing performance of incentive schemes. Although perhaps hard to implement in the radio context, fines are well known to be effective sanctions. We then analyze a jail-based mechanism and compare the performance to fines in Section IV. This is an extension of earlier work [32] where jails were used to discourage cheating if sensing were free. We conclude the paper in Section V by looking at the possibilities of aligning the incentives of primary and cognitive users for the purposes of cooperative system design.

## II. BASIC MODEL SETUP

We begin this investigation of deterrent mechanisms with a very simple toy model. We assume time is slotted, and the primary has a temporal behavior characterized by a probability  $p$  of turning off when it is currently on, and probability  $q$  of turning on when it is currently off. We assume this behavior is mostly memoryless, so the primary operation can be captured in a simple two-state Markov chain.

<sup>6</sup>The degrees of freedom we are concerned with here are bandwidth and time. If you have less bandwidth to send a fixed amount of data, the transfer will take more time. So, restricting one of these factors while requiring the other to remain constant restricts the radio's ability to transmit information.

<sup>7</sup>A similar model will apply to interaction within a real-time market system. For a discussion of this, see [31]

The cognitive user, in response to the primary activity, must choose optimal parameters in order to maximize its overall utility. We assume for simplicity that the cognitive user is risk-neutral, and therefore will only care about its expected utility over time. So, for now, the important characteristic of the primary can be captured by the overall probability that the primary is transmitting,  $P_{tx} = q/(q+p)$ .

The cognitive user can control its amount of sensing to optimize its utility. We assume that the cognitive user is employing a simple energy detector. So, the cognitive user takes measurements of the frequencies it would like to occupy and computes the test statistic [6]:

$$T(Y) = \frac{1}{N} \sum_{n=1}^N |Y[n]|^2 \quad (1)$$

where  $T(Y)$  is the test statistic, and  $Y[n]$  is the measurement of the channel at time  $n$ . We assume that the primary is transmitting at power  $P$  and experiences noise at the cognitive receiver  $\sigma^2$ . We are interested in the probabilities of error, which are the probability of false alarm (thinking the primary is there when it is not, quantified by the probability that the test statistic is greater than a threshold  $\gamma$  given that the primary is not active) and misdetection (thinking the primary is not active when it is in fact transmitting, quantified by a similar probability) [33]:

$$P_{fa} = P(T(Y) > \gamma | H_0) \quad (2)$$

$$P_{md} = P(T(Y) < \gamma | H_1) \quad (3)$$

where  $\gamma$  is a threshold parameter that can be chosen to optimize the detector,  $H_0$  denotes that the primary is not active, and  $H_1$  denotes that the primary is active. Using the Central Limit Theorem to approximate the distribution of the test statistic when the primary is present and when the primary is not, we can approximate the probabilities of false alarm and misdetection:

$$P_{fa} \approx Q\left(\frac{\gamma - \sigma^2}{\sqrt{2/N}\sigma^2}\right) \quad (4)$$

$$P_{md} \approx 1 - Q\left(\frac{\gamma - (P + \sigma^2)}{\sqrt{2/N}(P + \sigma^2)}\right) \quad (5)$$

where  $Q(x)$  is the probability that a standard normal Gaussian random variable is greater than  $x$ . If we vary the threshold,  $\gamma$ , from  $-\infty$  to  $\infty$ , then  $P_{fa}$  sweeps from 1 to 0, while  $P_{md}$  sweeps from 0 to 1 along a curve called the ROC (receiver operating characteristic). This ROC is a function of the SNR ( $10 \log_{10}(P/\sigma^2)$ , the ratio of the strength of the signal transmitted by the primary and the noise experienced at the cognitive user receiver in dB), and the amount of time spent sensing,  $N$ . Varying  $N$ , and fixing SNR at 0.2dB, we achieve the ROC curves shown in Fig. 1. As  $N$  increases, the receiver can simultaneously lower its  $P_{fa}$  and  $P_{md}$ . Holding  $N$  fixed and changing the SNR will also result in different ROC curves: as the SNR increases, the curves move toward the origin. We will consider what happens when SNR is varied in Section V.

For each of the primary's time steps, the cognitive user is assumed to have  $C = 1000$  steps of its own to spend. It must

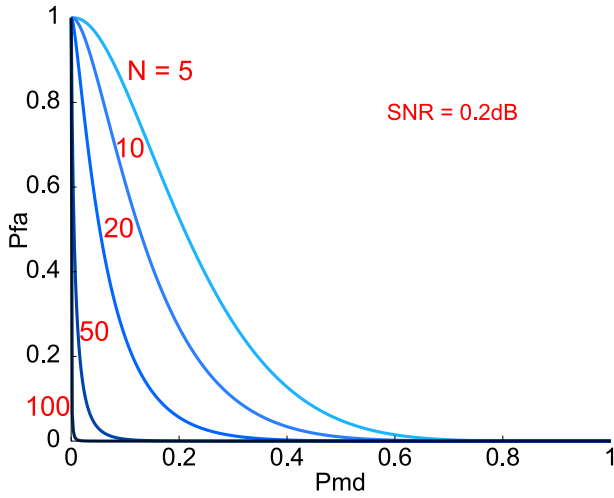


Fig. 1. The ROC curve for an energy detector, at different values of  $N$ . Similar curves can be obtained by changing the value of the SNR. However, the two cases are philosophically very different as raising  $N$  directly lowers utility for the cognitive user.

choose  $N$ , the number of steps to spend sensing, and therefore also  $C - N$ , the number of steps left over for transmission.<sup>8</sup> We will assume that the utility of the cognitive user is the total time it is able to transmit, so sensing takes up valuable time.<sup>9</sup> Using the parameters defined thus far, we can establish a base utility  $U_B$ , that does not yet include sanction terms:

$$U_B = (P_{tx}P_{md} + (1 - P_{tx})(1 - P_{fa})) \frac{C - N}{C} \quad (6)$$

The cognitive user can transmit if the primary is on and it misdetects, or if the primary is off and the cognitive user does not false-alarm. However, it can only transmit when it is not sensing ( $(C - N)/C$ ). Notice that without sanction, the cognitive user can maximize this utility by never sensing ( $N = 0$ ) and being oblivious to the primary's possible presence ( $P_{fa} = 0, P_{md} = 1$ ). So, without sanction, the rational cognitive user *always* causes harm to a transmitting primary.

In order to encourage cognitive users to sense, a regulator must be able to catch cheaters and issue sanctions. We assume that a method of determining identity, like that in [14] is in place. Then, we can apply parameters  $P_{catch}$  and  $P_{wrong}$ , which are the probabilities of catching a cognitive user currently causing harm, and the probability of issuing a sanction to an innocent cognitive user, respectively.<sup>10</sup> We assume that sanctions can only be applied when the primary

<sup>8</sup> $N$  could also be thought of as an amount of time spent getting permission to use the band. The cycle repeats every primary time step.

<sup>9</sup>We are assuming that the secondary gets utility even if there is some pollution caused by the primary's presence in the band. This is actually not a bad assumption. When the cognitive user is degree-of-freedom limited instead of interference limited [34], it does not mind having pollution from a primary in the band. See [35] for a discussion of the differences between the two kinds of users and their preferences regarding reclaiming unused spectrum in polluted bands.

<sup>10</sup>Throughout, we are assuming strict liability (for a discussion of strict liability, see Posner [24]). So, if a cognitive user is caught, it will be assessed a sanction without any opportunity to defend itself or demonstrate that it exercised a reasonable amount of care.

user is active, so a cognitive user cannot be punished if there is no-one around to harm.

With this basic model, we are ready to explore the first approach to deterrence: applying fines to misbehaving radios.

### III. DETERRING CHEATING WITH FINES

Whether monetary or otherwise, we think of a fine as a direct decrease in utility. In order for the fine to have desirable marginal deterrence capabilities, the amount of the fine must rise with the amount of interference the cognitive user is causing the primary. So, we define a fine  $F$  that is issued whenever the cognitive user is caught interfering (innocently or not). Thus, the overall utility function for a cognitive user under the fine system is:

$$U_F = U_B - FP_{tx}(P_{md}P_{catch} + (1 - P_{md})P_{wrong}) \quad (7)$$

The goal of the cognitive user, then, is to find the  $N$  and  $P_{md}$  which maximize its utility:

$$\{N^*, P_{md}^*\} = \underset{N, P_{md}}{\operatorname{argmax}} U_F \quad (8)$$

We use  $P_{md}$  as a proxy for the real optimization variable,  $\gamma$ ,

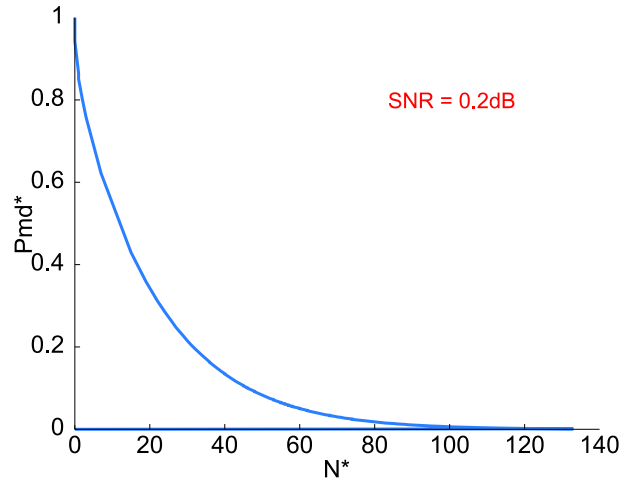


Fig. 2. The tradeoff curve between the optimal number of samples and the optimal probability of missed detection. Regardless of enforcement parameters, this curve seems to stay the same.

the threshold on the energy detector. Choosing how much to misdetect is important because it is controlling how much a cognitive user is choosing to cheat given the information that the sensing provides. It may be that sensing is so costly that you want to sense very little. But the fine is also costly, so you want to lower your  $P_{md}$  to actively avoid the fine (this will have the effect of raising your probability of false alarm,  $P_{fa}$ ).

As the fine varies, the optimal values of  $P_{md}$  and  $N$  vary as well, along the curve shown in Fig. 2. There are two things to notice about this plot: first, fines simply work. No matter what protection is desired for the primary (where protection is measured by  $P_{md}$ , the probability of misdetection), setting an optimal fine will correctly direct the cognitive user's behavior. Second, the other parameters of the problem,  $P_{tx}, P_{catch}, P_{wrong}$ , influence what the fine must be to achieve

a particular place on this curve, but do not affect the optimal curve itself.<sup>11</sup> We will revisit this second point later when we talk about the influence of SNR.

Our definition of when the cognitive user can be caught for causing harm suggests that the correct metric of good cognitive behavior should be the average amount of time it interferes with an active primary:

$$I_F = \frac{P_{tx} P_{md} \frac{C-N}{C}}{P_{tx}} \quad (9)$$

We are interested in what effect raising the fine has on the utility for the cognitive user and the average interference caused to the primary user. This is shown in Fig. 3

Notice that there are two very distinct regions: if the fine is too low, the cognitive user has no incentive to sense ( $N^*/C = 0$ ), or to set its  $P_{md}$  below one. As the fine is raised, the sanction becomes more expensive than sensing, and the cognitive user will transition to a sensing regime. After this transition, the average interference to the primary drops off quickly. When implementing a fine system, the average interference to the primary is almost equal to the misdetection probability for the cognitive user. It differs only in the amount of time spent sensing,  $N/C$ . The remaining time  $(C-N)/C$  is essentially one whenever  $P_{md}$  is high, and insignificant when  $P_{md}$  is low. The utility for the cognitive user holds fairly constant over the sensing regime. Because of the threshold behavior, we say the fine is effective if the secondary is in the sensing regime.

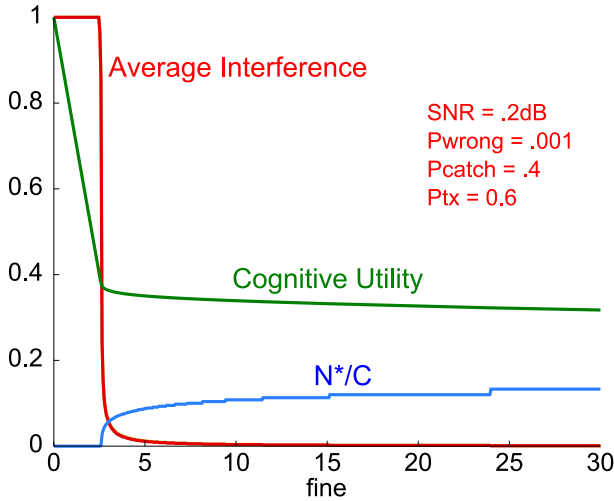


Fig. 3. The effect on primary and secondary parameters of interest varying the amount of the fine. Notice that there are two distinct regions: one where the cognitive radio will not sense, and one where it will. When it is sensing, it is causing very little interference to the primary

The fine required to push the cognitive user to sense depends on the transmission characteristics of the primary, as shown in Fig. 4. When the primary is rarely active, it is difficult to force them to sense. This conforms with our intuition as when the primary is almost never there, the fine will only be incurred

very infrequently. So, the fine must be much higher to get the cognitive users to sense. Unfortunately, this means that if a regulator wants to protect a range of primary behavior, including a primary that is almost never active, it must set an extremely large fine thereby hurting even the honest cognitive users.

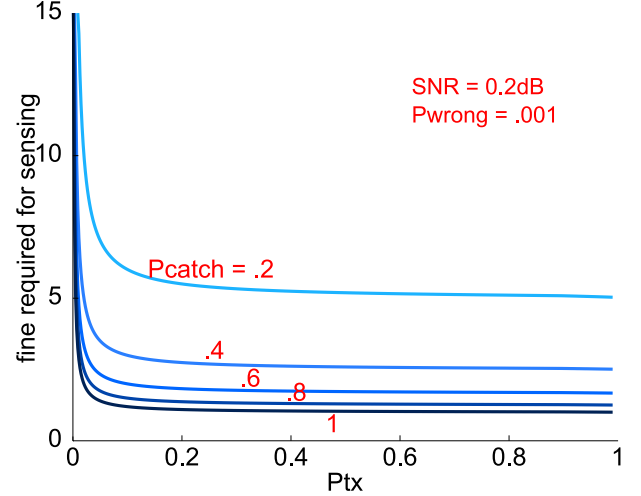


Fig. 4. The required level of fine to incentivize the secondary to sense vs probability the primary is transmitting for different values of  $P_{catch}$ .

The required fine also depends non-trivially on the difference between  $P_{wrong}$  and  $P_{catch}$ . Fig. 5 shows the tradeoff of the optimal number of samples and  $P_{wrong}$  for different values of  $P_{catch}$ , and a fixed fine. As  $P_{wrong}$  gets close to  $P_{catch}$ , the incentive to sense decreases, and eventually, the cognitive user will no longer sense. The intuition here follows very nicely that developed in [36]: if there is a non-trivial probability of being wrongfully convicted, this probability decreases the threat of punishment through your own actions since you will likely incur the punishment anyway. Therefore, you might as well cheat.

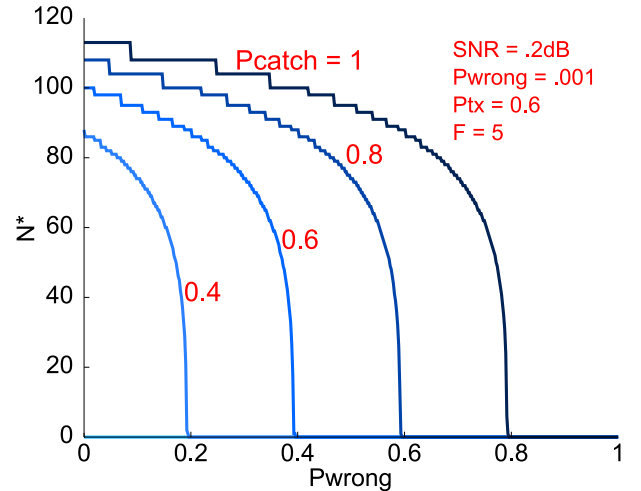


Fig. 5. The optimal number of samples  $N^*$  vs  $P_{wrong}$  for different values of  $P_{catch}$ . As  $P_{wrong}$  gets closer to  $P_{catch}$ , for a given level of fine, the cognitive radio will lose incentive to sense at all.

<sup>11</sup>The authors have not thus far been able to prove that this is true in general. But simulations give evidence that suggest it holds in several test cases beyond the energy detector.

To gain further insight into the required fine, we plot in

Fig. 6 the overhead the cognitive user is experiencing because of the fine. We define overhead as the utility the cognitive user would have gotten at the same level of sensing without the fine:

$$O_F = \frac{U_B - U_F}{U_B} \quad (10)$$

Notice that the cognitive user will start sensing when it is losing a significant fraction of its utility to the fine. This figure also shows the overhead of the cognitive user relative to what utility it could get if it knew exactly when the primary was transmitting. As soon as this overhead is non-negative (the cognitive user is no longer gaining from always cheating), it will sense. Pushing the cognitive user to a moderate level of sensing produces the best results in terms of overhead. So, overall, the cognitive user has the best overhead performance with a moderate level of sensing, and the primary user has a low average interference at the same time. Therefore, this is the likely preferred operating point.

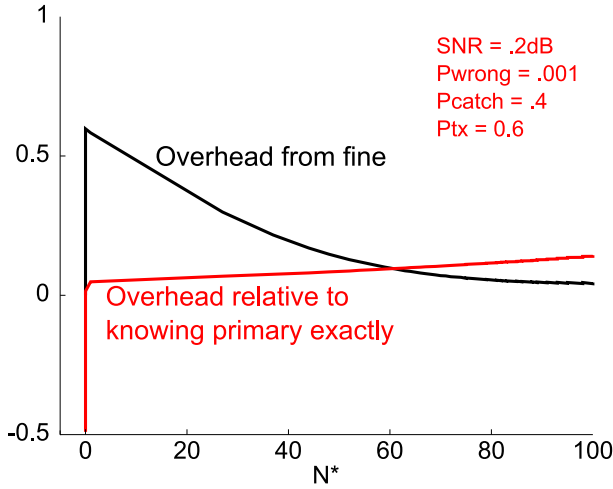


Fig. 6. The overhead incurred for the fine system, defined as the percentage of the base utility that a cognitive radio would have to pay in order for  $N^*$  time steps to be in its best interest. The cognitive user will sense when the fine becomes too expensive. Also, the overhead of a fine system relative to what the cognitive user would get if it knew exactly when the primary was transmitting. The cognitive user will sense as soon as cheating does not bring any benefit (negative overhead).

Employing sanctions in the form of fines is obviously effective – the fines can be used to deter cheating for any desired level of protection for the primary. Moreover, fines operate with a threshold behavior: there is a critical fine at which the cognitive users will start sensing. Above this critical point, the fine has very little further effect. So it operates much like a reserve price.

However, these fines are effective because they scale with the utility of the cognitive user. This is a well known effect in literature on fines as deterrents in the criminal justice system. When the wealth of the criminal increases, so must the fine in order to be an effective deterrent [26].

Unfortunately, wealth is not always observable, or may be very costly to calculate. Polinsky studies both of these cases in a pair of papers, [37], [38]. If information about the wealth is not available at all, the optimal approach is to offer the criminal a choice of sanctions between a high fine or a low

fine plus jail sentence (this assumes two wealth classes, but can be generalized to more). If the wealth is simply hard to observe, there must be a self-reporting with a chance of audit and a higher imposed fine if the criminal is found to have lied about his wealth.

Can these results be applied to our case? A radio could be forced to admit how much it has transmitted, which is our proxy for wealth. But how does this map to a monetary fine? Even if the fine is determined by a percentage of total utility, the mapping would be difficult. The natural step would be to assess the fine at the service-provider level as a percentage of the total earnings, but this does not take into account utility that is not reflected in profits. It also does not allow individual devices to be punished for bad behavior, so the deterrent effect would be reduced by moving the punishment further from the crime.

Assessing fines at the service-provider level would have the further complication of allowing the fine to be treated as a price of operation (see [39] for a case study of this effect, and see [24] for the general theory) instead of allowing it to have a deterrent effect. If fines are actually as difficult to implement correctly as they seem, perhaps a better option is to work with a different form of sanction. Traditionally, the most natural other sanction is imprisonment through a jail system [23]. We will investigate the radio equivalent in the next section.

#### IV. DETERRING CHEATING WITH JAILS

Fines are convenient to analyze, but hard to implement. So, in this section we will extend the previous analysis to one of the universal currencies of radios: degrees of freedom. Note that we could alternatively have chosen to analyze energy. If a radio were energy-constrained, like a sensor network running on limited battery power, losing energy is painful. One could imagine that an effective sanction for this type of device could be forcing the radio “sing in the corner.” A dedicated frequency would be blocked off for punishment, and radios that were caught would have to transmit a beacon on this frequency, thus wasting precious energy. We will not consider the energy limited case, but investigation is warranted to better understand how to universally deter bad behavior by radios without knowing what they specifically care about.<sup>12</sup>

We will consider here the most obvious kind of cognitive user: one that is constrained by how many degrees of freedom it is able to obtain. Therefore, it can be deterred from cheating with a simple punitive action modeled after jails in the human realm – a spectrum jail in which offenders cannot transmit. This means that if a cognitive device is caught transmitting when the primary is active, it must turn off its radio for a period of time.

Notice that while in the human realm jails are a burden to society because they are costly to implement, in the radio domain, they are not. The cost for implementing is the cost of certifying that a radio can obey a “go to jail” command

<sup>12</sup>This singing in the corner could even serve a useful purpose – if all radios undergoing punishment had to sense constantly and use the band to give a truthful account of the primary activity they observe, the singing could be effectively public service. We will not consider this case here but it would be an interesting extension.

by turning itself off, and the cost of sending that “go to jail” command. A jail in spectrum does not require maintenance. There is an *opportunity cost* to society overall because the jailed radio is unable to perform any socially-valuable function. However, for a radio that is not sensing, the harm to the primary may be greater than the good it produces by using the band. We can interpret this as incapacitation.

All other necessary costs, like those associated with catching offenders are incurred in the fines case as well, and are a necessary cost of implementing a policing scheme for spectrum. Therefore, unlike traditional reasoning for the human realm [23], jails may be the preferable option for sanctions in spectrum.

### A. Modeling jails

In [32], we developed a simple model for a jail system in which sensing is free, that employs the same basic parameters used for fines in the previous section. We extend this same model here to include the cost of sensing and choice of how much to sense and how much to misdetect.

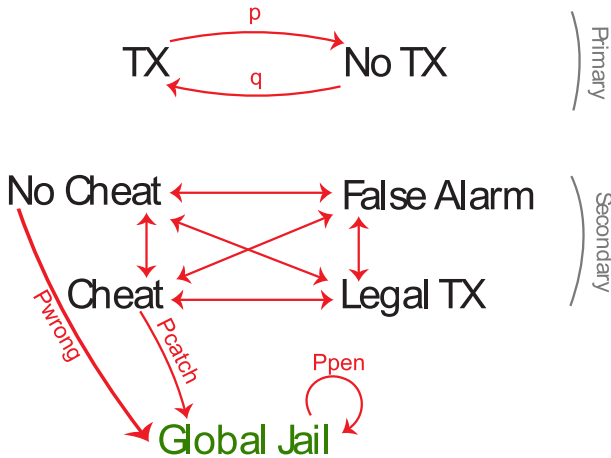


Fig. 7. The Markov model which governs the behavior of the primary and secondary users under a jail system.

The basic model is shown in Fig. 7. The primary has the same model as before, and the cognitive user responds in the same way. But now if the cognitive user is caught interfering with the primary (correctly or erroneously), it is sent to jail, where it must turn off its radios for a length of time that is geometric with expected value  $1/(1 - P_{pen})$ . To keep the comparison to the fines case simple for now, we will assume that we are dealing with an iid primary. This means that  $p+q = 1$  so  $P_{tx} = q$ . In this case, we can calculate a simple utility function that has a similar form to the one we used before:

$$\tilde{U}_J = \pi_{NJ} U_B \frac{C - N}{C} \quad (11)$$

The cognitive user gets utility whenever it is not in jail (which happens with probability  $\pi_{NJ}$ ) and transmitting. The  $\pi$  notation indicates the stationary probability of being in a particular set of states in the corresponding Markov chain. This is equivalent to the long-term average amount of time that you spend in the particular state. The probability that it

is not in jail is simple in this case, because we can collapse the model above into a two state Markov chain with states “in jail” and “not in jail”. Then, the probability of not being in jail is:

$$\pi_{NJ} = \frac{1 - P_{pen}}{1 - P_{pen} + P_{tx}(P_{md}P_{catch} + (1 - P_{md})P_{wrong})} \quad (12)$$

It turns out that this definition is insufficient to incentivize sensing in all cases. Consider when the primary is almost always active ( $P_{tx} \approx 1$ ). The cognitive user will get no utility if it just sits idle, whereas it gets some positive utility by cheating, going to jail, then cheating again when it is released. No length of jail sentence would be sufficient to force it to sense.

We encountered this same problem in the development of a model in which sensing was free in [32], and solved it using a “home band”<sup>13</sup> of value  $\beta$  and a global jail. The home band is a clean piece of spectrum where the cognitive user holds primary rights and so always has access. In order for the cognitive user to expand into other bands, it must stake this home band against improper use. If the cognitive user is caught cheating in another band, it must turn off its radios in *all* bands, including its home band, so it loses utility when sitting in jail (whereas before the utility of being in jail was simply zero). The new utility function including the home band is:

$$U_J = \pi_{NJ}(P_{tx}P_{md} + (1 - P_{tx})(1 - P_{fa})) \frac{C - N}{C} - \beta\pi_J \quad (13)$$

where  $\pi_J = 1 - \pi_{NJ}$  is the stationary probability of sitting in jail. The cognitive user is trying to maximize this utility over its available parameters. It computes:

$$\{N^*, P_{md}^*\} = \operatorname{argmax}_{N, P_{md}} U_J \quad (14)$$

Using the same parameters of interest as we did in the fines case, we can look at the qualitative behavior for different levels of jail punishment (which we will parameterize by  $\beta/(1 - P_{pen})$  for reasons justified later). This behavior is shown in Fig. 8.

The same two regions exist: a region in which the cognitive user will never sense, and a region in which the cognitive user will sense. There is the same effect of the average interference to the primary dropping after the the cognitive user starts sensing. However, as in the human realm, even if the cognitive user cheats at every possible opportunity, the jails have an incapacitating effect (for a general overview of this effect, see [25]). The cognitive users will certainly cause harm to the primary if they are allowed to continue to transmit, so by sending them to jail, the primary users are better protected than they would have been had just a fine been imposed.

This figure also shows directly the overhead caused by the jail system: the cognitive user will begin to sense when it is spending too much time in jail relative to the time it would have spent sensing. After this, the overhead incurred by the jail

<sup>13</sup>The homeband for cognitive users can be thought of as being similar to the bond a bonded contractor has to post. It is something of value that you must stake against the project failing.  $\beta$  could also be thought of as an extra punishment for jail. Perhaps, again, while a cognitive device is in jail, it must sing in another dedicated band. Then, while sitting in jail it is losing energy.

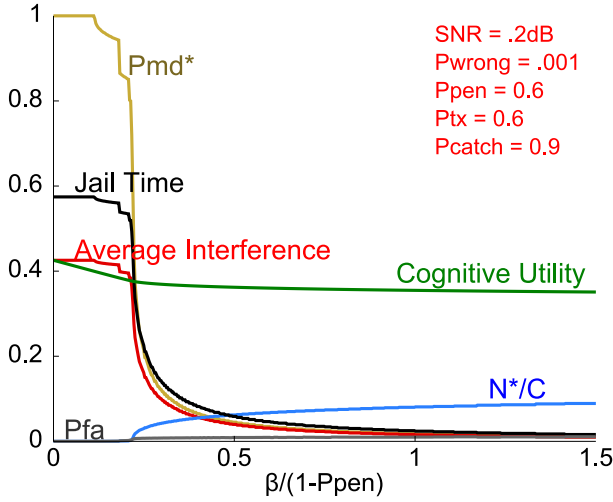


Fig. 8. Primary and cognitive user parameters of interest in the jail system. Again, we have the behavior of very sharp regions where the cognitive user will sense or not. There is also an incapacitation effect to jail in the non-sensing region.

system itself is relatively low. The sensing regime, despite its slightly lower overall utility, may be the preferable option for many kinds of cognitive users. If an application is sensitive to delay, like streaming video, having to wait in jail is a very costly possibility. The delay is more important even than the overall amount of time transmitting. Therefore, a delay-sensitive device will likely prefer to switch to the sensing regime even for lower values of  $\beta/(1 - P_{pen})$ .

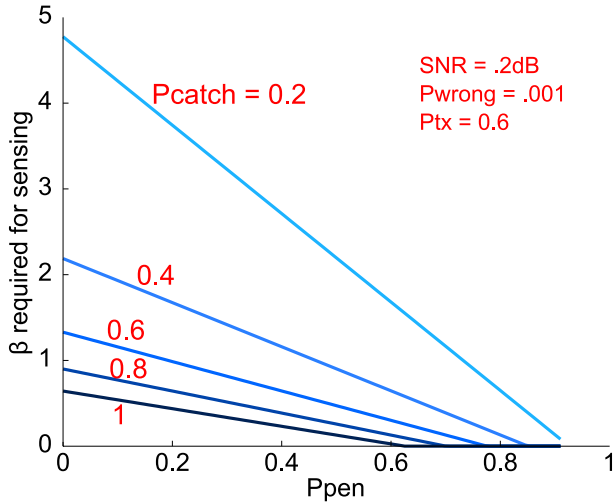


Fig. 9. This plot shows the tradeoff between the minimum  $P_{pen}$  and  $\beta$  required to deter cognitive users from cheating.

With the same two sensing and non-sensing regions, we are again concerned with the level of punishment that produces the transition. We begin with  $\beta$  and  $P_{pen}$ , the parameters unique to the jail system. Fig. 9 shows the minimum  $\beta$  required for sensing to occur given different values of  $P_{pen}$  and fixed other parameters. Notice that  $\beta$  and  $P_{pen}$  trade off linearly, so both affect the secondary response in the same way. This makes sense intuitively: every time the cognitive user is sent to jail, it experiences an amount of punishment with average value

$\beta/(1 - P_{pen})$  (the pain of jail times the average length of stay). We should be able to shorten a jail sentence while making it more painful and keep the same deterrent effect. We will use this  $\beta/(1 - P_{pen})$  parameterization throughout the paper to describe the amount punishment associated with jail.

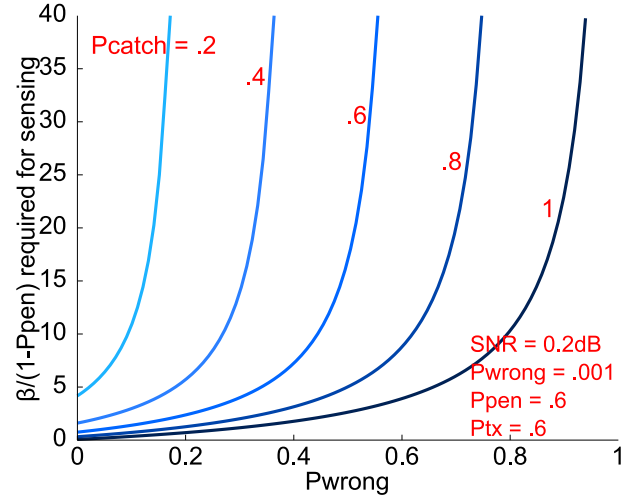


Fig. 10. The tradeoff between  $P_{wrong}$  and  $\beta$  to deter cheating.

Looking at the relationship between  $P_{catch}$ ,  $P_{wrong}$ , and  $\beta/(1 - P_{pen})$  in Fig. 10, we see the same effect as we saw with fines: as  $P_{catch}$  approaches  $P_{wrong}$ , the jail sanction required to incentivize sensing goes to infinity.

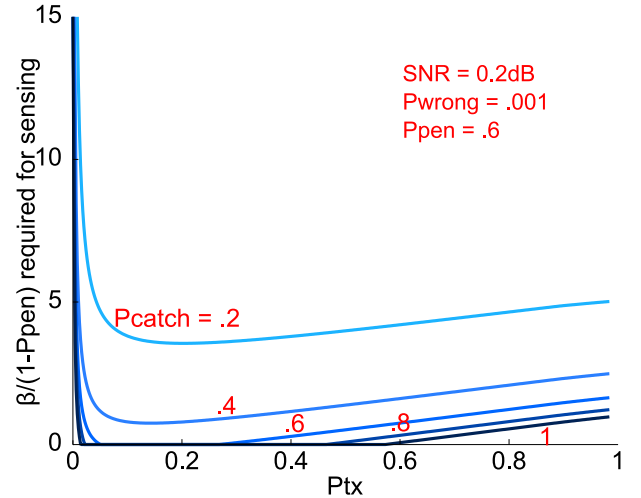


Fig. 11. The optimal amount of sensing for different levels of primary activity employing a jail system. The difficult regions to deter cheating in this case are when  $P_{tx}$  is either very low or very high.

We finally look at the required sanction against the probability the primary is transmitting, as we did for the fines case. This is shown in Fig. 11. We see an additional problem area: when the primary is almost never around, there is very little incentive to sense because sensing costs all the time, while jail costs only infrequently. The primary always being present is also problematic because the cognitive user can gain utility by bouncing in and out of jail. Use of the homeband can correct both problem areas, but at the expense of honest cognitive users in situations of  $P_{tx}$  in the middle range.

### B. Equivalence of fines and jails




Fig. 12. Fines and their corresponding equivalent  $\beta/(1 - P_{pen})$ . When the  $P_{md}$  and  $N$  are evaluated at their optimal values for particular levels of fines, the fine and equivalent jail punishment enjoy a linear relationship.

We have seen thus far that the qualitative behavior of jails is equivalent to the behavior we observed with fines. Indeed, even the tradeoff curve between  $N^*$  and  $P_{md}^*$  in Fig. 2 is the same for jails. The jail parameters simply trace out a different section of this line. This all suggests that there is a mapping between a fine and a jail sanction. This mapping is:

$$F_{eq} = \frac{U_B + \beta}{1 - P_{pen} + P_{tx}(P_{md}P_{catch} + (1 - P_{md})P_{wrong})} \quad (15)$$

Notice that the mapping depends on the secondary chosen parameters  $P_{md}$  and  $N$ . So, from this equation alone determining the optimal jail sanction from the optimal fine is difficult. However, when this equation is evaluated at  $P_{md}^*$  and  $N^*$  for each particular value of fine, we get the equivalence curve shown in Fig. 12. The two sanctions enjoy an almost linear relationship. So, in the iid primary case, a jail sanction can be as effective as a fine in deterring cheating and indeed the jail parameters can be set by appropriately scaling the optimal fine. Jail even has the added benefit of simplicity of implementation and immediate incapacitation of non-compliant secondaries.

### C. Non-iid Primaries

Although we have shown fines and jail sanctions to be in some sense equivalent in the iid primary case, primaries are very rarely iid. A primary is much more likely to stay in the on position for a while before switching off. We refer to this as the “stickiness” of the primary: its tendency to stay in one place for a longer time given a certain overall probability of being on. We can quantify the stickiness of the primary by looking at the second eigenvalue of the primary transition matrix. A second eigenvalue of zero produces the iid case we have been dealing with thus far. A second eigenvalue closer to one will produce a primary more likely to stick in particular states before transitioning.

When the primary does not have an iid behavior, the utility function the cognitive user is is more complicated because it

must take into account the fact that jail sentences may occur while the primary is sticking in the on or off position. The utility function for the secondary is now:

$$U_{SJ} = \pi_{(NJ,on)}P_{md} + \pi_{(NJ,off)}(1 - P_{fa}) - \beta\pi_J \quad (16)$$

The difference between this and Eq. (11) is that the probability that the primary is on and the cognitive user is not in jail,  $\pi_{(NJ,on)} \neq \pi_{NJ}P_{tx}$  as it did in the iid case. Likewise, the probability the primary is not on when the cognitive user is not in jail is  $\pi_{(NJ,off)} \neq \pi_{NJ}(1 - P_{tx})$ .

Note that the utility function for fines in the sticky-primary case remains unchanged: the fine is assumed to be assessed on the total average utility, not as a part of the temporal operation of the cognitive user. So, the effect of sticky primaries shows up only when dealing with a jail system.




Fig. 13. Tradeoff between optimal secondary parameter values for different levels of primary stickiness. Notice that the base tradeoff remains unchanged. The only difference is what parts of the curve are reachable.




Fig. 14. Parameter values for different levels of stickiness and different jail sanctions. Notice that the primary difference is a translation in the point at which the cognitive user switches to the sensing regime.

Fig. 13 shows the optimal tradeoff between  $N^*$  and  $P_{md}^*$  for different levels of stickiness for the primary. Note that the curve itself is the same as it has been throughout. Nothing changes the actual optimal tradeoff. However, different levels of stickiness do affect the effectiveness of the enforcement

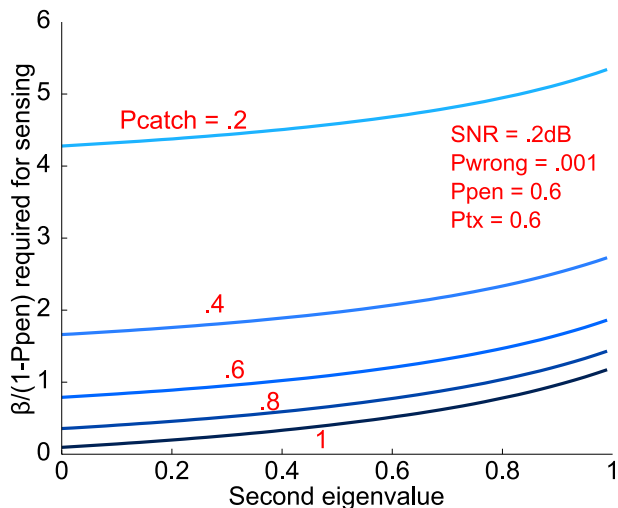


Fig. 15. The  $\beta/(1 - P_{pen})$  required for the secondary to sense for different levels of primary stickiness (different values for the second eigenvalue of the primary transition matrix).

parameters. This can be seen by the jump between not sensing and a particular level of sensing. Once the cognitive user is forced into the sensing regime, it does not necessarily (at least to the level of precision we have calculated here) begin sensing with one time step. It jumps to a higher level of sensing. The different jumps in the curve indicate that the different cases are in fact changing the relative effectiveness of the enforcement parameters.

This same effect can be seen in Figs. 14 and 15. In the first, we compare the qualitative effect of an increasing jail sanction on the parameters of interest. Notice that the two regions remain. The main difference is the value of jail sanction that manages to switch the secondary user into a sensing regime. The location of this point is shown against the second eigenvalue of the primary in Fig. 15. As the primary gets more sticky (second eigenvalue going to one), the jail sanction is less effective. This is because if the primary user is likely to stay transmitting for a long time, going to jail represents a drain on utility (through the  $\beta$  factor), but it does not represent a loss of opportunity to legally transmit. As the primary gets more sticky, the probability that there was no legal transmission opportunity during a jail sentence grows. So, jail is painful only from a  $\beta$  perspective, not because of an opportunity loss. As the second eigenvalue goes to one, the primary looks the same as a primary that is always present or always absent, and we end up with the problem areas on  $P_{tx}$  we observed in Fig. 11. Therefore, when considering how to design a jail system for radios, a regulator must take into account the characteristics of the primary, not just in how much time it intends to spend transmitting, but also in how long it intends to transmit at a time.

## V. THE EFFECT OF COOPERATION

We have thus far considered SNR to be a fixed parameter inherent of the system at runtime. However, “runtime” actually consists of two distinct components. At engineering design cycle time, systems are designed and tested. At the actual

runtime, systems are deployed and used. Although the policing system must operate at actual runtime, the systems could be designed to favor more peaceful coexistence.

SNR is a parameter that can capture design-time considerations. While SNR does represent the actual signal to noise ratio observed at runtime, cooperative strategies like those in [8] can change the *effective* SNR experienced. For example, with the “OR” rule, a group of cognitive users claims a primary is present if *any* of the users observe the primary. Any individual radio independently sees the true SNR, but the network cares about the *highest* SNR received, which raises the effective SNR. Likewise, if the primary were to transmit a pilot signal (which greatly increases the probability they will be detected [6]), this too could be interpreted as a larger effective SNR. So, we are interested here in whether these cooperative strategies are in the best interest of either primary or cognitive users.

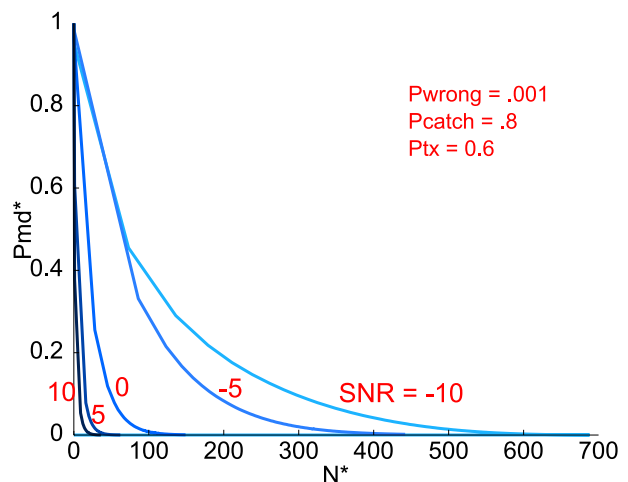


Fig. 16. The effect of SNR on the optimal tradeoff between the number of samples and the probability of missed detection.

First, note Fig. 16. SNR is a parameter of the ROC curve for the energy detector our cognitive devices are assumed to be using. Therefore, as the SNR changes, the relationship between the sensing time  $N$  and the probability of missed detection  $P_{md}$  must also change. Fig. 16 shows the curve for the optimal  $N^*$  versus  $P_{md}$  for different values of SNR. As SNR grows, the optimal curve moves toward the origin – as SNR grows, it becomes easier to achieve a low  $P_{md}$  from fewer samples.

The effect of changing the SNR on the parameters of interest for primary and cognitive users is shown in Fig. 17. As we vary SNR, we see the same qualitative sensing regions: for very low SNR and the same level of sanction, the cognitive user will never sense, while for high SNR, the cognitive user will sense. Notice also the large jump in the number of samples taken close to the transition into the sensing regime. Smaller SNRs require more samples to produce a favorable ROC. Therefore, the number of samples taken is highest right after transitioning to sensing and will drop as the SNR continues to grow.

The most notable aspect of this plot is cognitive utility on either side of the transition. When the cognitive user is sensing, its utility is significantly higher because it is not spending a large portion of its time in jail. Also, the cognitive user requires fewer samples at higher SNR, so sensing is not

causing overhead either. As in all plots before, the average interference to the primary drops significantly in the sensing regime. So, *both* users prefer to operate in the sensing regime. This alignment of desires indicates that primaries may have incentive to transmit a pilot to help the cognitive users sense. Likewise, the cognitive user may have incentive to use cooperative schemes to help raise the SNR. Of course, this depends on the relative cost of employing such strategies to the benefits that will be obtained, but the opportunity exists.




Fig. 17. These sets of parameters look at the reactions of cognitive users as we vary the operating SNR. The same sense/not-sense boundary occurs, and will have implications for possible cooperation between primaries and secondaries.




Fig. 18. We keep these parameters against SNR, but now look at the difference in the parameters when  $P_{wrong}$  is no longer very small. This may have implications for self-policing among cognitive users.

One final note on cooperation concerns  $P_{wrong}$ . As we saw before,  $P_{catch} - P_{wrong}$  is the quantity that determines the required sanction for sensing. So, if  $P_{wrong}$  is large for a given  $P_{catch}$ , the punishment must be very severe and wrongful convictions will hurt a lot.  $P_{wrong}$  has essentially three sources: noise, primaries, and other cognitive users. The noise is random noise happening to look like significant interference,

leading to a conviction. If the methods of encoding identity are done well, this effect should be small.

The primary can be a major source of wrongful convictions by simply crying wolf: if the primary claims interference often enough, there is a higher chance of noise being mistaken for a particular user's identity. In our current model, the primary is essentially a powerless player, and so this effect will not show up here.<sup>14</sup>

The final source of wrongful conviction comes from other cognitive users. If one of the other users is constantly cheating, eventually your radio will be convicted along with them. The intuition comes from the identity codes described in [14]. These are codes stamped into the time series a user is transmitting. A radio can be identified simply by observing the on-off pattern of transmission without the need to understand what the signal means. In order to support large numbers of active users, random groups of users must share the same identity code for a period of time. If someone in your group is caught cheating, the whole group will be punished. Therefore, the wrongful conviction of every cognitive device will be raised if even one other device is constantly cheating.

Consider Fig. 18. This is the same setup as Fig. 17, but now we have a significantly higher probability of wrongful conviction ( $P_{wrong} = 0.2$ ). The sensing region is still well defined, albeit occurring at a higher SNR. The primary still observes a drop in the probability that they will be interfered with. However, the cognitive user gains almost no utility from switching to the sensing regime, whereas before they got a considerable bump. The cognitive user is not spending as much time in jail, but it is also not getting any utility from cheating. So, overall, its utility does not rise much because it does not get enough time back from jail to offset the loss in cheating utility. From this plot, it is obvious that it is of interest for the cognitive users to lower  $P_{wrong}$ . It may be possible that cognitive users even have incentive to punish cheaters on their own outside of the confines of the jail system.

This external punishment might look like shunning in the human realm. In [40], it is shown that for transmitting messages over long distances, it is optimal to pass that message along a series of relay nodes instead of trying to transmit it directly to the receiver. Therefore, any shunned radio whose messages are not relayed by its neighbors will waste considerable power trying to transmit its messages by itself. As long as cognitive users could distinguish who was cheating (perhaps another reason to make radios "sing in the corner"?) they could employ shunning as a deterrent to keep the other cognitive users honest.

## VI. CONCLUDING REMARKS

In an earlier paper, [32], we introduced a model to understand whether a jail-based sanction were sufficient to convince a cognitive user to respect sharing rules given it knew the actions of a primary. The intent was to explore a possible alternative to certification-only regulation that would have a

<sup>14</sup>It would conceivably be possible to create a system of incentives, like making it costly to accuse a cognitive user of interference, that could lower the probability of wrongful conviction due to primaries crying wolf.

more acceptable overhead. In this paper we see that the same jail mechanism is sufficient to incentivize the cognitive user to sense for primary presence, even though sensing is costly. Furthermore, the benefits of sensing show a threshold behavior: once the secondary begins sensing, stronger sanctions derive little extra benefit for the primary. In the course of this investigation, we made a number of interesting experimental observations. We saw that in the space of possible sensing times and mis-detection probabilities, the optimal choices lie along a single curve. This phenomenon deserves further study because if it is true in general beyond the energy detector considered here, there may be important policy implications.

We also saw that the preferences of the primary and the secondary align regarding the perceived SNR at the secondary: both do better when that SNR is higher and the primary is more easily detectable. This points to possible cooperative strategies, but a more thorough investigation is warranted to treat the primary as an active player and analyze equilibria.

Finally, we have looked at only one possible sanction. Because devices can also be energy-limited, using “singing in the corner” or public service transmission of sensing data may also be effective strategies. Ideally, an overhead analysis of these strategies would be performed to understand what is viable for future regulation of cognitive devices.

## REFERENCES

- [1] “Spectrum policy task force report,” Tech. Rep. 02-135, Federal Communications Commission, Nov. 2002.
- [2] M. A. McHenry and K. Steadman, “Spectrum occupancy measurements, location 1 of 6: Riverbend park, Great Falls, Virginia,” tech. rep., Shared Spectrum Company, 2005.
- [3] G. R. Faulhaber, “Wireless telecommunications: Spectrum as a critical resource,” *Southern California Law Review*, vol. 79, Mar. 2006.
- [4] “FCC Adopts Rules for Unlicensed Use of Television White Spaces.” [http://hraunfoss.fcc.gov/edocs\\_public/attachmatch/FCC-08-260A1.pdf](http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf).
- [5] F. Perich, “Policy-based network management for NeXt generation spectrum access control,” in *Proceedings of the Second IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, (Dublin, Ireland), Apr. 2007.
- [6] R. Tandra and A. Sahai, “SNR walls for signal detection,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, pp. 4–17, Feb. 2008.
- [7] R. Tandra, S. M. Mishra, and A. Sahai, “What is a spectrum hole and what does it take to recognize one?,” *Proceedings of the IEEE*, Jan. 2009.
- [8] S. Mishra, A. Sahai, and R. Broderon, “Cooperative sensing among cognitive radios,” *IEEE ICC*, pp. 1658 – 1663, June 2006.
- [9] D. Hatfield and P. Weiser, “Toward Property Rights in Spectrum: The Difficult Policy Choices Ahead,” *CATO Institute*, Aug. 2006.
- [10] R. H. Coase, “The Federal Communications Commission,” *Journal of Law and Economics*, vol. 2, pp. 1–40, Oct. 1959.
- [11] E. Goodman, “Spectrum Rights in the Telecom to Come,” *San Diego Law Review*, vol. 41, no. 269, 2004.
- [12] T. W. Hazlett, “Optimal Abolition of FCC Spectrum Allocation,” *Journal of Economic Perspectives*, vol. 22, no. 1, pp. 103–128, 2008.
- [13] P. J. Weiser and D. N. Hatfield, “Policing the Spectrum Commons,” *Fordham Law Review*, vol. 74, no. 2, pp. 663–694, 2005.
- [14] G. Atia, A. Sahai, and V. Saligrama, “Spectrum enforcement and liability assignment in cognitive radio systems,” in *Proceedings of the Third IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, (Chicago, IL), Oct. 2008.
- [15] Y. Benkler, “Overcoming Agoraphobia: Building the Commons of the Digitally Networked Environment,” *Harvard Journal of Law and Technology*, vol. 11, pp. 287–400, Winter 1998.
- [16] R. Etkin, A. Parekh, and D. Tse, “Spectrum sharing for unlicensed bands,” in *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, (Baltimore, MD), Nov. 2005.
- [17] G. R. Faulhaber, “Deploying Cognitive Radio: Economic, Legal, and Policy Issues,” *International Journal of Communication*, vol. 2, pp. 1114–1124, 2008.
- [18] K. Werbach, “Supercommons: Toward a Unified Theory of Wireless Communication,” *Texas Law Review*, vol. 82, pp. 863–973, March 2004.
- [19] L. Kaplow and S. Shavell, “Economic analysis of law,” *NBER Working Paper*, Feb. 1999.
- [20] C. Montesquieu, *The Spirit of the Laws*. University of California Press, Berkeley, 1977.
- [21] C. Beccaria, *An Essay on Crimes and Punishments*. W.C. Little, Albany, 1872.
- [22] J. Bentham, *An introduction to the principles of morals and legislation, in: The Utilitarians*. Anchor Books, Garden City, NY, 1973.
- [23] G. S. Becker, “Crime and Punishment: An Economic Approach,” *Journal of Political Economy*, vol. 76, 1968.
- [24] R. A. Posner, *Economic Analysis of Law*. Little, Brown, and Company, 3 ed., 1986.
- [25] A. M. Polinsky and S. Shavell, “The theory of public enforcement of law,” *Journal of Economic Literature*, vol. 38, pp. 45 – 76, Mar. 2000.
- [26] A. M. Polinsky and S. Shavell, “The economic theory of public enforcement of law,” *Journal of Economic Literature*, vol. 38, pp. 45 – 77, Mar. 2000.
- [27] H. Palaiyanur, K. A. Woyach, R. Tandra, and A. Sahai, “Spectrum zoning as robust optimization,” *Proceedings of the Fourth IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Apr. 2010.
- [28] A. Sen, “The impossibility of a Paretian liberal,” *Journal of Political Economy*, vol. 78, no. 1, pp. 152 – 157, 1970.
- [29] J. Rawls, *Justice as Fairness: A Restatement*. Belknap Press of Harvard University Press, 2nd ed., 2001.
- [30] M. Weiser, “The computer for the twenty-first century,” *Scientific American*, Sept. 1991.
- [31] A. Sahai, K. Woyach, G. Atia, and V. Saligrama, “A technical perspective on light-handed regulation for cognitive radios,” *IEEE Communications Magazine*, pp. 96 – 102, Mar. 2009.
- [32] K. Woyach, A. Sahai, G. Atia, and V. Saligrama, “Crime and punishment for cognitive radios,” *Allerton*, 2008.
- [33] H. L. V. Trees, *Detection, estimation, and modulation theory*. Wiley-Interscience, 2001.
- [34] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, May 2005.
- [35] A. Sahai, K. Woyach, K. Harrison, H. Palaiyanur, and R. Tandra, “Towards a “theory of spectrum zoning”,” *Allerton*, Oct. 2009.
- [36] I. P. L. Png, “Optimal subsidies and damages in the presence of judicial error,” *International Review of Law and Economics*, vol. 6, pp. 101 – 105, June 1986.
- [37] A. M. Polinsky, “The optimal use of fines and imprisonment when wealth is unobservable,” *Journal of Public Economics*, vol. 90, no. 4-5, pp. 823 – 835, 2006.
- [38] A. M. Polinsky, “Optimal fines and auditing when wealth is costly to observe,” *International Review of Law and Economics*, vol. 26, no. 323, 2006.
- [39] U. Gneezy and A. Rustichini, “A fine is a price,” *Journal of Legal Studies*, vol. 29, Jan. 2009.
- [40] P. Gupta and P. Kumar, “The capacity of wireless networks,” *IEEE Transactions on Information Theory*, vol. 46, Mar. 2000.