

A toy-model for the regulation of cognitive radios

Kristen Woyach and Anant Sahai
 Wireless Foundations
 Department of EECS
 University of California at Berkeley
 Email: {kwoyach, sahai}@eecs.berkeley.edu

Abstract—With the development of frequency agile radios comes a new option for using wireless spectrum: cognitive radio. Under this scheme a hierarchical approach is employed in which legacy “primary” users may use the band as they wish, while “secondary” users may take advantage of any unused spectrum. However, as of yet the question of enforcement in cognitive radio is unaddressed, especially in the context of cooperative spectrum sensing. Unfortunately, it cannot be assumed that secondaries will not cheat as there is no physical limitation to stop them. Therefore, enforcement mechanisms and models to evaluate these mechanisms are required. This paper considers the problem of enforcement in terms of incentives available to encourage secondaries to play by the rules. It develops a toy game-theoretic model to capture the dynamics of the primary and secondary users and evaluates what this model can tell us regarding the overall tradeoff between enforcement mechanisms, cheating behavior and channel utilization. In particular we find that in order to guarantee that there is no incentive to cheat, it is not enough to punish cheaters simply by denying them access to the band. Also, we find that it is important to catch the individual cheater: no amount of punishment can incentivize playing by the rules when collective punishment is employed.

I. INTRODUCTION

As wireless technology becomes more ubiquitous, spectrum availability becomes an increasingly important concern. However, despite full allocation by the FCC, empirical measurements suggest spectrum is in fact vastly underutilized [1]–[3]. The real question becomes why is this precious resource being wasted and how can we correct the situation? We discuss some related perspectives on this problem in Section II and then decide that any solution to this problem must involve regulation that involves some form of explicit or implicit policing at runtime. This paper focuses on the case of opportunistic spectrum users that want to use a band in which a primary is only intermittently active.

Ideally, we want to be able to answer questions such as the following:

- How should regulation be partitioned between device certification, standards, and incentives?
- How easily is this enforced?
- At what level is the regulation performed (FCC, standards body, individual run-time decisions)?
- Does this regulation scheme encourage innovation and technological improvement?

In the companion paper, [4], the enforcement of cognitive radios is considered to understand the overhead required and the tradeoffs involved to detect interference and catch culprits

reliably. However, that story must be complemented with a quantitative argument that bounds what is realistic from an incentive perspective. Otherwise, the parameters of any code could never be set. This paper begins the process of understanding these incentives and the tradeoffs in enforcement and utilization. This is done by defining a toy model to capture the enforcement dynamics in the cognitive-radio scenario.

Obvious tradeoffs exist between utilization, the enforcement mechanism, and the cheating behavior of the secondary users. Consider the extreme cases: if no enforcement were employed, there is no accountability and therefore no incentive to follow the prescribed hierarchy. Secondaries would choose to transmit with no regard for interference to the primary. Although utilization would be high, the primaries would be severely affected, negating the original hierarchy concept entirely. On the other hand, if the punishment for cheating is that the secondary user is forever denied access to the band, the penalty is so severe that the secondaries would have to be very conservative in their use in order to avoid it. They must be so conservative, in fact, that the overall utilization would be essentially the same as if there were no cognitive users at all. In this paper, we seek to understand this tradeoff.

To that end, in Section III we identify the important dynamics and present a simple model to capture them. Sections IV and V investigate what this model can tell us about the tradeoffs of interest, and Section VI presents some concluding remarks on where we go from here.

II. BACKGROUND

The problem of spectrum underuse naturally lies with the current allocation system and so regardless of technological improvements, regulatory questions are unavoidable. When the FCC was designed, “wireless” use was synonymous with “broadcast” and interference was something to be feared [5]. The allocation system was designed on scales appropriate for broadcast. The allocation time was considered on the order of ten years: the order of device lifetimes. Geographic scales were considered in the hundreds of square kilometers, or the order of a broadcast service area. At the time, this model worked well because the number of players was relatively small and interference was the primary concern.

However, this original model is no longer appropriate because it produces a mismatch between the scales of allocation and scales of use. Many of the current developments in wireless networks are geared toward smaller, more decentralized

concepts which simply do not operate at the large scales assumed for broadcast systems. The effect is the existence of “holes” [6] in spectrum usage. In time these appear as legacy networks move to packetized, service upon demand schemes that do not constantly require spectrum use. Time holes also exist when legacy systems no longer have the ability or incentive to operate and simply sit idle on their allocated band. In space, these holes occur as the legacy systems operate in a service area that does not cover the entire allocated area.

If the FCC had the information and the speed to adapt appropriately, it could dynamically allocate spectrum to the best users in real time. Unfortunately, as a centralized regulation authority that defines even the location and height of transmit antennas for every piece of spectrum in the United States, the FCC is not physically capable of having all the necessary information and operating at the appropriate speed.

It seems that utilization would improve if more flexibility were allowed, but this is only half the story. When spectrum is heavily regulated, the rules of use are clear, and anyone not following the rules is easily caught and punished. The current FCC actually represents an extreme position in which a high degree of regulation produces very simple enforcement. If you know where and how all the transmit antennas are operating, it is very simple to weed out illegal parties with a directional antenna. At the other extreme, consider a case in which there is no regulation. [7] suggests through game-theoretic arguments that the system would self-regulate so that everyone has some operational ability. The outcome would look qualitatively similar to the internet: the QOS would be best effort. It would be only as good as the worst player is willing to accept and it would be very difficult to effectively contain malicious users.

The optimal solution in trying to rethink the rules, then, must lie somewhere in the middle. Ideally we want a “light-handed” regulation in which operational rules and means of enforcement are defined to maintain peaceful coexistence. The rules must also be general enough to provide flexibility and encourage innovation. When considering spectrum use, regulating either only *a priori* with equipment or only at run time with usage rules and policing seems insufficient. However, the correct mix of regulation scales is yet unclear.

Solutions to the allocation problem have been proposed, but differ widely in their approaches. The first is spectrum privatization, introduced by Coase in [5] and elaborated on by de Vany in [8]. The idea here is to assign property rights to spectrum much like they exist for land. This would include a time, geographic area, and spectrum band. Correct usage is defined as not exceeding a power threshold at the boundary of one’s region. Beyond this stipulation, users are permitted to use the band however they wish. With property rights, the spectrum could be sold, divided, merged, and otherwise renegotiated between individuals instead of by a central authority. Instead of a regulatory body making decisions, market forces determine the overall utilization. The argument is that this approach would converge to the most efficient use of spectrum because those most able to effectively use the spectrum are precisely those who are most able to purchase the property rights. Note that the scales of allocation would begin on the

order of the original FCC allocation scheme, but these too would presumably converge to the correct scales of use that most effectively utilize available spectrum.

The ability to enforce in this scheme comes down to distinguishing when and who is at fault. But, this is not necessarily simple. As Hatfield points out in [9], the wireless medium is variable and propagation depends on many factors including even the current state of the ionosphere. So, it is not sufficient to simply determine maximum power levels *a priori* and assume they will always work. From the regulatory perspective, it is much easier to agree on a propagation model and then have users prove their models are in compliance. But this too raises difficult questions. Even if a good model could be developed, it is not obvious that it would even be possible to determine who is at fault. It is also not clear how to treat interference claims: should interference be considered as trespass or nuisance [10]? On the one hand, trespass seems appropriate because interference is one user overstepping their property and infringing on someone else. This is nice because litigation here is relatively straightforward. However, this treatment can encourage “spectrum troll” [9] behavior in which owners can simply sit idle on their spectrum, wait for someone to transmit too loudly, and sue them. Perhaps a more appropriate concept is nuisance, which evaluates the degree to which one is affecting the other and can weigh the relative utility of the users to determine who is at fault. Unfortunately, although this allows flexibility which will help maximize utilization, actually resolving nuisance cases is subjective and difficult.

A second school of thought is characterized by Benkler [11] which treats spectrum as a commons. Here, the concept is to place the regulation burden solely on the equipment instead of distributing it between equipment and use, so that users are naturally able to coexist when using the spectrum in whatever manner they choose. The market, then, is pushed from the spectrum itself to the equipment and will favor those technologies more able to coexist effectively. Here we do not have to worry about usage scales as operators may use their devices however they wish. But it is not clear how long device certification should last or if the standards could evolve effectively. The flexibility and adaptability of this route is obvious: users are free to innovate with no more risk than the initial hardware investment. Enforcement is concerned with original certification of devices and finding malfunctioning units. Again, this becomes a technical question: can the equipment be certified to guarantee that coexistence is possible and how do we recognize offending parties in this sea of users who are relatively unconstrained?

Cognitive radio emerges as a third regulatory option built like a hybrid of the previous two. It seeks to solve the allocation problem through a hierarchy of operation regimes. The legacy “primary” systems, with rights to the spectrum either through allocation or market transaction, may continue to use the spectrum and operate on the scale defined by the regulation scheme. However, “secondary” users may then operate at a much smaller scale, observing holes in spectrum usage and filling them as long as their operation does not interfere with the primary users. Here, again, the potential

for higher utilization and innovation is obvious. We reclaim unused portions of spectrum, increasing utilization. Innovation is encouraged as primaries improve their *services* from competition with secondaries, and the secondaries have the freedom to experiment so long as they conform to the hierarchy. Many papers exist on what the secondary must be able to do, technically, to operate under these constraints (see for example [12]), but the question of enforcement is largely ignored. It is generally assumed that the secondaries will implicitly follow the rules.

This is unfortunately not a fair assumption. There are no physical device constraints that limit cheating, so if one wants to regulate cheating behavior directly with standards it would require sifting through thousands of lines of code for every certified device. As this is not feasible, we must therefore consider enforcement from a perspective of how to incentivize secondaries to not cheat. Again, a game-theory approach seems appropriate to capture these dynamics.

With this in mind, we can question how to meet the goal of light handed regulation. Unfortunately, the answer is not obvious because all of these approaches treat the problem in very different ways. Although all present obvious ways of increasing utilization and encouraging innovation, they also leave large technical questions unanswered, primarily in the realm of enforcement. But, even if the mechanisms for enforcement were clear, we have no metrics to define what is good nor an understanding of the fundamental tradeoffs involved.

It could be argued that defining metrics is unnecessary as we could simply design an experiment where each proposal was pitted against the others and market forces would decide the most efficient solution. However, with unused spectrum already at a premium, such an experiment may be infeasible and certainly could not be performed blindly. Even if we could perform such an experiment, what works in one case is not necessarily applicable to another band and usage scenario. What are really needed are models that capture the dynamics of interest and give a qualitative and quantitative sense of the tradeoffs involved. With these tradeoffs in hand, appropriate metrics can be defined to compare different ideas against the ideal.

This motivates the need for the idealized toy model in the next section.

III. MODEL

The goal of our proposed model is to capture the basic dynamics of the primary and secondary users to get a handle on some of the tradeoffs involved. Specifically, a good model of this situation should help us:

- Understand what needs to be regulated and what does not
- Understand the burden on the regulator
- Compare different penalties
- Understand the effect of different technologies

We are concerned here primarily with parameter choices and responses, so we will make simplifying assumptions that remove complications. First, we assume that the primary and secondary users exist at a single point in space. This allows the

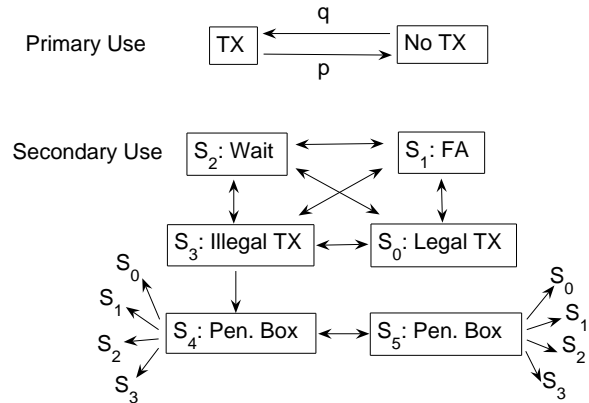


Fig. 1. Markov chain for modeling secondary cheating dynamics in cognitive radio system. The primary use is modeled as the two-state chain at the top, which determines the horizontal movement through the secondary chain. The secondary chain responds to the primary's use by transmitting or waiting on either side. Illegal use of the band prompts a penalty, which is sitting in the penalty box for a specified amount of time.

suppression of any propagation effects so that interference can be treated as a binary quantity: if both primary and secondary are transmitting, it qualifies as interference. We also suppress overlaps that would normally occur between the time the primary turns on and the time the secondary notices. Likewise, the time lag between the primary turning off and the secondary noticing is also suppressed.

We will assume three possible players in this game: the primary determines channel use and cares about whether it is experiencing interference. However, the primary usage is assumed fixed, and so the primary is a silent, oblivious player. The secondary is assumed to always have something to transmit, and so it cares only about maximizing its transmit time. The regulator is able to determine the level of enforcement and cares about overall channel utilization.

The dynamics of the game are modeled by movement in the Markov chain depicted in Fig. 1. The primary, which can be transmitting or not, moves according to fixed parameters p and q . When the primary is not transmitting, the secondary can either be legally transmitting (S_0) or registering a false alarm (S_1). When the primary is transmitting, the secondary may be waiting as it should (S_2), or illegally transmitting (S_3). The penalty boxes (S_4 and S_5) represent the punishment for cheating: if a secondary is caught cheating, it must sit in the penalty box for a specified amount of time.

Horizontal movement through the secondary chain is determined by the primary's channel usage, so it depends on p and q . Vertical movement is determined by the regulator and secondary setting individual parameter values to maximize respective utility functions. The regulator's goal is to maximize utilization while suppressing cheating. Therefore, it uses the utility function:

$$U_R = \max_{P_{catch}, P_{pen}, \beta} \pi_0 + \pi_2 \quad (1)$$

where π_i is the stationary probability of being in state S_i . The secondary cares only about being able to transmit, and so its




Fig. 2. p_{FA} vs p_{MD} tradeoff curves for an energy detector. Bold lines are at SNR = -6dB with number of samples $N = 400$ and $N = 1000$

utility function is

$$U_S = \max_{p_{cheat}, p_{FA}} \pi_0 + \pi_3 - \beta(\pi_4 + \pi_5) \quad (2)$$

where β is a variable factor that determines how much the penalty boxes hurt the secondary beyond simply not being allowed to transmit (this case corresponds to $\beta = 0$). The physical meaning and implications of this factor will be discussed later in Section IV-C.

To maximize their utility functions, the users set parameters. The regulator controls parameters related to catching cheaters and punishing them. This includes p_{catch} , the probability that a cheating secondary will be caught, p_{pen} , which determines the average length of stay in the penalty boxes, and β , the variable extra punishment for sitting in the penalty boxes. p_{catch} and β include external constraints which will be discussed later in Section IV-C, while p_{pen} can be freely set.

The secondary controls parameters related to cheating. This includes p_{cheat} , the probability of choosing to cheat given you know the primary is transmitting. This is the secondary's free parameter with no external constraint. The secondary also controls p_{FA} and p_{MD} , the probabilities of false alarm and missed detection, respectively. These two parameters are not independent. Indeed, we can assume that they lie on a curve determined by the secondary's sensing mechanism. For all simulations here, the detector is assumed to be an energy detector which has the approximate test statistic [12]

$$T(Y)|\mathcal{H}_0 \sim \mathcal{N}\left(\sigma^2, \frac{1}{N}2\sigma^4\right) \quad (3)$$

$$T(Y)|\mathcal{H}_1 \sim \mathcal{N}\left(P + \sigma^2, \frac{1}{N}2(P + \sigma^2)^2\right) \quad (4)$$

This produces the family of curves shown in Fig. 2 where the shape of the curve is determined by the energy detector itself and the slope of the curve is determined by the SNR and number of samples N . The middle bold curve is the nominal one used for most of the simulations. When a better curve is required for comparison, the other bold curve is used.




Fig. 3. p_{cheat} against channel use, characterized by the stationary probability the primary is transmitting and p , the probability the primary will move from transmitting to not transmitting in the next step. In this low penalty regime where $p_{catch} = 0.8$, $p_{pen} = 0.6$, and $\beta = 0$, the cheating behavior always takes a binary form, with more cheating occurring when the channel is less available.

To determine vertical state transitions, the parameters work together as follows:

- To enter S_0 , the secondary must not experience a false alarm ($1 - p_{FA}$)
- A false alarm (p_{FA}) determines transitions into S_1 .
- To enter S_2 , the secondary must know the primary is transmitting and choose not to cheat ($(1 - p_{cheat})(1 - p_{MD})$).
- To enter S_3 , the secondary must know the primary is on but cheat anyway ($(1 - p_{MD})p_{cheat}$) or not know the primary is talking (p_{MD}).
- To be sent to the penalty box, the secondary must already be in S_3 and get caught (p_{catch}).
- Once in the penalty box, the secondary can leave if it is not forced to stay ($1 - p_{pen}$).

IV. TRADEOFFS OF INTEREST

We explore how this model answers the desired questions by first qualitatively understanding the behavior of the secondary with respect to the channel usage of the primary, and how it affects the overall utilization. We then take a quantitative look at the effect of the regulator's enforcement parameters, and the effect of changing secondary technology.

A. Effect of primary channel use

In considering usage of the primary, we want to get a sense of whether "one size fits all": can we set the enforcement parameters once for any channel, or do they need to adapt based on channel usage? To answer this question, we fix the enforcement parameters and let the secondary adjust its parameters to maximize its utility function. This results in two modes of operation.

In the first, shown in Fig. 3, the penalty is not bad compared to the opportunity to claim extra transmission time. So, the p_{FA} is pushed as low as the simulation will allow and p_{cheat} is a binary quantity which is determined by the enforcement



(a) p_{FA} (b) p_{MD}

Fig. 4. p_{FA} and p_{MD} against channel use when $p_{catch} = 0.8$, $p_{pen} = 0.6$, and $\beta = 3$. When the penalty is high enough, the secondary begins to avoid it by raising its own p_{FA} .

parameters and channel use. In this figure, the z-axis is the probability that the secondary will cheat, either implicitly with p_{MD} or explicitly with p_{cheat} . Cheating is visualized for different values of p , the probability that the primary transmitting in the current state will be not transmitting in the next, and the steady state probability that the primary is transmitting ($q/(p+q)$). The steady state probability is used as a natural general parameter because even if the primary had a more complicated usage pattern, it still has some sort of duty cycle. The dependence on the steady state probability of transmitting is intuitively pleasing as we would expect the secondary to cheat more when the channel is less available. p is used to directly see the effect of choosing between cheating now with possible penalty or legally transmitting in the near future. In all plots, the allowed range of the probability the primary is transmitting and p is outlined at the bottom for reference. Note that this region is a result of defining the model as a discrete Markov chain. If the model were continuous, the region would not be restricted.

Fig. 4 shows the second mode of operation in which the penalty is so high that the secondary actively avoids getting caught. In this case, p_{cheat} is always zero, and p_{FA} rises to




Fig. 5. Channel utilization if only the primary is present.

push p_{MD} as low as possible. Notice again the dependence on the probability that the primary is transmitting. When this value is high, there is greater opportunity to get caught, so the measures taken to reduce this probability are greater, i.e. p_{FA} is higher.

B. Overall channel utilization

From a regulatory perspective, the metric of interest is the overall utilization of the channel: we want to know whether the current enforcement scheme allows the secondary to effectively fill spectrum holes. So here we will discuss the effect of different modes of operation on the overall utilization.

All the cases that follow can be compared to Fig. 5, which shows the utilization when only the primary is present. It is equivalent to the steady state probability of the primary transmitting. For all other utilization figures, two definitions are used: The first is total utilization, defined as

$$U_{total} = \pi_0 + \pi_2 + \pi_3 + \pi_4 \quad (5)$$

which is the percentage of time anyone is using the channel. The second definition, utility with collisions, is defined as

$$U_{collide} = \pi_0 + \pi_2 + \pi_4 \quad (6)$$

which does not count interference as utilization.

Fig. 6 shows the utilization for the cheating behavior observed in Fig. 3. Note in this figure that when the secondary is always cheating, the utilization follows more closely the curve in Fig. 5. When the secondary is not cheating, it more effectively fills spectrum holes. This effect is magnified when considering collisions; here the cheating drags the utilization down significantly when the primary is usually transmitting.

Fig. 7 shows the overall utilization when the secondary is avoiding punishment, with p_{FA} as shown in Fig. 4. When the primary is always transmitting, it dominates utilization even when discounting collisions. When the primary is rarely talking, the utilization is good because the secondary has a low p_{FA} in this region. The middle, however, shows relatively poor utilization as the secondary is still actively avoiding getting caught even though the primary is not fully using the band.

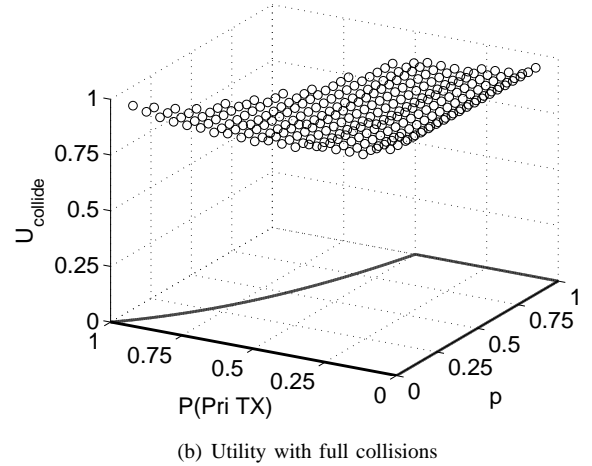
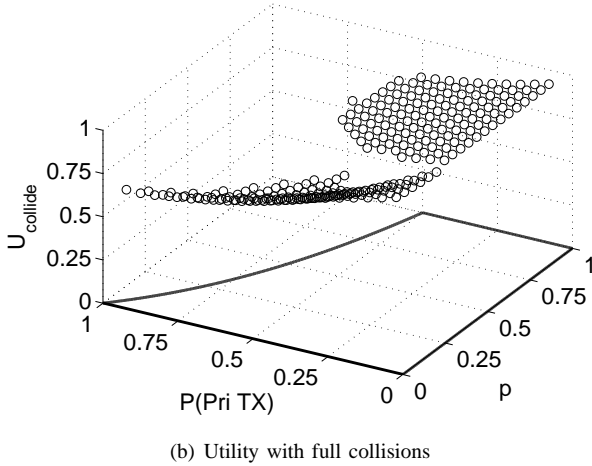
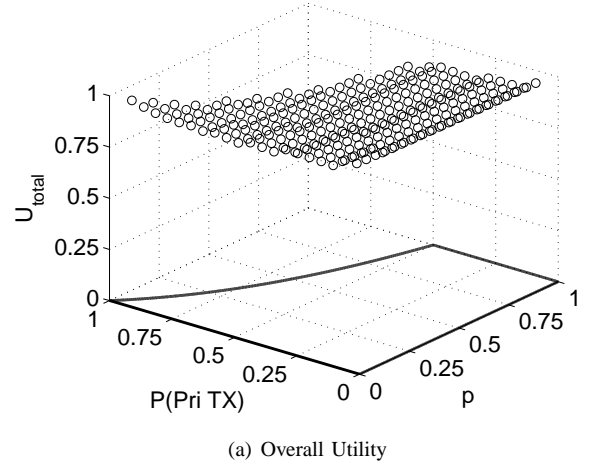



Fig. 6. Overall utility (percentage of time either the primary or secondary is using the channel, without collision) and utility with full collision (if both primary and secondary are transmitting, not counted toward utility) in a low penalty regime with $p_{catch} = 0.8$, $p_{pen} = 0.6$, and $\beta = 0$.

Fig. 7. Effect of conservative FA on utility $p_{catch} = 0.8$, $p_{pen} = 0.6$, and $\beta = 3$

Fig. 8 shows the utilization in between the last two cases: the secondary is not cheating, but it is not actively avoiding the penalty either. Notice that qualitatively this curve looks much like that in Fig. 7, but it gets better utilization in the middle, where before the secondary had an inflated p_{FA} . This is the sweet spot for which any regulation should be aiming because it is the point of maximum utilization.

Fig. 9 shows the case where p_{MD} is held fixed at 0.05. The utilization is approximately equal to our incentive scheme, so we lose little with light-handed regulation. The difference then is in simplicity of enforcement. [13] and [14] suggest that to achieve a good p_{MD} , cognitive radio systems will have to employ some kind of cooperation. This means that in order to regulate a fixed p_{MD} , one would have to certify software by sifting through thousands of lines of code. With an incentive system, however, one simply has to correctly define the enforcement parameters.

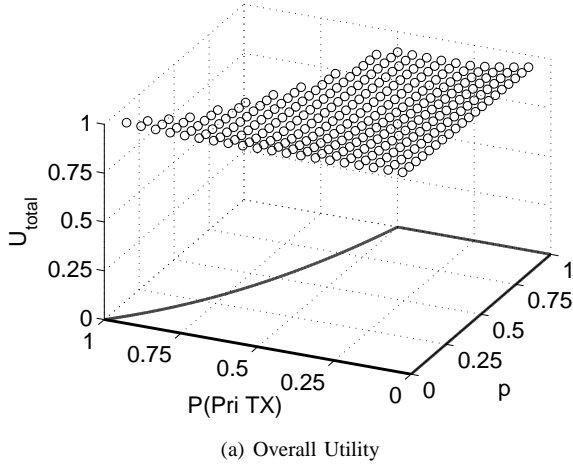
C. The effect of enforcement parameters

In our model, the regulator has access to three parameters that determine enforcement: p_{catch} , β , and p_{pen} . We want to


get a sense of what these physically mean and what effect each has on the secondary cheating behavior. We would also like to get a sense of how much of a burden enforcement is on the regulator. To do this, we will isolate the effects of different parameters and give physical interpretations of the results.

p_{catch} captures the enforcement mechanism itself by representing the chance a penalty will be employed. The probability of catching cheaters depends on how the catching is done and the deployment of catching nodes. Because secondary cheating always exhibits a binary behavior, we can evaluate the effect of p_{catch} by tracing how the boundary between always and never cheating moves when p_{pen} and β are held fixed (Fig. 10). p_{catch} determines the location of the boundary with respect to the probability of the primary transmitting, but it does not affect the shape much.


β is a factor that, in conjunction with p_{pen} , affects the type and severity of the punishment. Physically, β could represent a number of things. It can be thought of as an extra fine imposed when in the penalty box. It can also be used to represent a cost of missed opportunity if we were to consider a multiband scenario in which being in the penalty box in one band denies use in *any* band. As such, β could be a highly constrained




(a) Overall Utility



(a) Overall Utility



(b) Utility with full collision



(b) Utility with full collision

Fig. 8. Utilization with no cheating or extra false alarms $p_{\text{catch}} = 0.8$, $p_{\text{pen}} = 0.6$, and $\beta = 0.5$

Fig. 9. Utilization when p_{MD} is fixed at 0.05. Overall utilization is not much better than with a purely incentivized scheme.

parameter. If it is considered in the multiband sense, it is affected by the usage of the primaries in other bands as well as how interested particular secondaries are in using those bands. So, β could even be different for different secondaries. If a value is assigned, the effect of this parameter on behavior is to give a relative weight between the profit gained by using the band and the penalty incurred from using it improperly. In regimes where the secondary is cheating, the effect of β is the shift shown in Fig. 11 which looks much like that of p_{catch} .

p_{pen} determines the length of stay in the penalty box, effectively determining the dependence of cheating on p . All examples thus far have p_{pen} at a moderate value. To get a better sense of its effect on behavior, consider the extremes. When $p_{\text{pen}} = 0$ as in Fig. 12, the secondary will stay in the penalty box for only one time step. So, whether it cheats is based on two things: the cost of cheating and whether the next time step will be an opportunity for legal transmission. The first is determined by p_{catch} and β as before, but the second is determined solely by the transition probability p and does not depend on channel usage. Therefore, the boundary between cheating and not in this case is a function only of p .

When p_{pen} is very high, as in Fig. 13, the secondary will

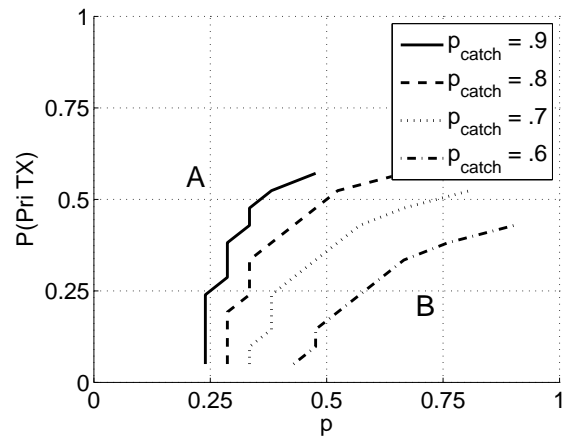


Fig. 10. While keeping $p_{\text{pen}} = 0.6$ and $\beta = 0$ fixed, the boundary between cheating (region A) and not cheating (region B) migrates with different values of p_{catch} , but the shape does not change.




Fig. 11. While keeping $p_{catch} = 0.4$ and $p_{pen} = 0.3$ fixed, the boundary between cheating (region A) and not cheating (region B) migrates with different values of β , but the shape does not change.




Fig. 12. When p_{pen} is very low ($p_{pen} = 0$ here), with $p_{catch} = 0.8$ and $\beta = 1$, the boundary between cheating and not cheating is determined only by the chance of legally transmitting in the next time slot, p .




Fig. 13. When p_{pen} is very high ($p_{pen} = 0.99$ here), with $p_{catch} = 0.02$ and $\beta = 0$, the choice to cheat or not cheat depends on the long term average channel availability, so it is determined almost entirely by the steady state probability of the primary transmitting

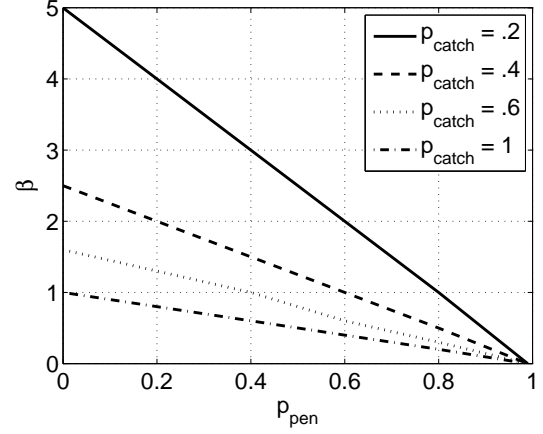


Fig. 14. Bounding β values vs. p_{pen} where no more cheating occurs. Plotted for different values of p_{cheat}

have to spend a long time in the penalty box, so the time to return to legal transmission does not matter as much as a long term average of channel availability. Therefore in this case, the dependence is almost entirely on whether the primary is transmitting and not on p .

In a real system, these enforcement parameters can be chosen in several ways. If the primary usage is known a priori, they could be minimally set so that there is no incentive for the secondary to cheat. However, as usage may be less predictable or unknown, we could instead use bounding values that ensure the secondary has no incentive to cheat regardless of the primary usage. These bounding values were found empirically and are plotted in Fig. 14. As p_{catch} is the technology-limited parameter, it is used to determine which β - p_{pen} tradeoff curve is needed. These lines follow exactly an intuitive mathematical argument: if we define the total penalty as

$$K = \frac{\beta}{1 - p_{pen}} \quad (7)$$

which is the extra β factor times the expected amount of time in the penalty box, then the secondary would be tempted to cheat if

$$K p_{catch} < 1 \quad (8)$$

$$\frac{\beta}{1 - p_{pen}} p_{catch} < 1 \quad (9)$$

or the secondary is tempted to cheat if the cost of the penalty is less than the utility gained by transmitting in that time slot. Equality in (8) determines the boundary between cheating and not; the boundaries found in Fig. 14 closely follow this expression. Note that this assumes a particular kind of traffic and QoS for the secondary. It will fit well the case when the secondary does not have time constraints, so the choice to cheat is based solely on the relative utility. If the secondary has time-sensitive data, however, the model will change as the secondary will have greater incentive to cheat to maintain a constant connection.

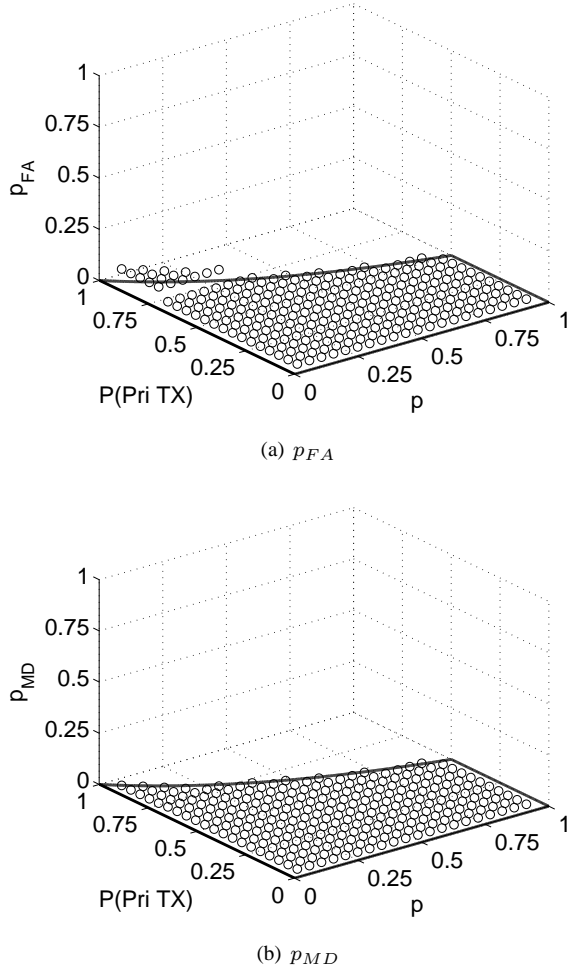


Fig. 15. With $p_{catch} = 0.8$, $p_{pen} = 0.6$, and $\beta = 3$, a better p_{FA} - p_{MD} curve means that the secondary does not have to increase its p_{FA} (and decrease p_{MD}) much to avoid penalty.

D. The effect of secondary technology

We can model the effect of better secondary detection technology by simply replacing the p_{FA} - p_{MD} curve with a better one from Fig. 2. Under the same conditions that produced the levels of p_{FA} and p_{MD} in Fig. 4, the better tradeoff curve produces Fig. 15. Notice that while the original tradeoff curve requires the secondary to actively avoid the penalty, this one does not because the p_{MD} is already low enough so that the penalty is not incurred too often. With the improvement in p_{FA} comes an improvement in overall utility, shown in Fig. 16.

Note that another tradeoff is suppressed here. The better p_{FA} - p_{MD} tradeoff can also be achieved by simply sensing the channel longer. However, the better curve achieved by sensing longer must be traded off against the opportunity to use the available channel and may not always be worthwhile. So, it is possible that the detector curve should be a function of the channel usage as well which is not yet captured here.

If the secondary has a perfect detector, regardless of the strength of the penalty, p_{FA} and p_{MD} will never increase above zero, and the utility will be maximized as soon as

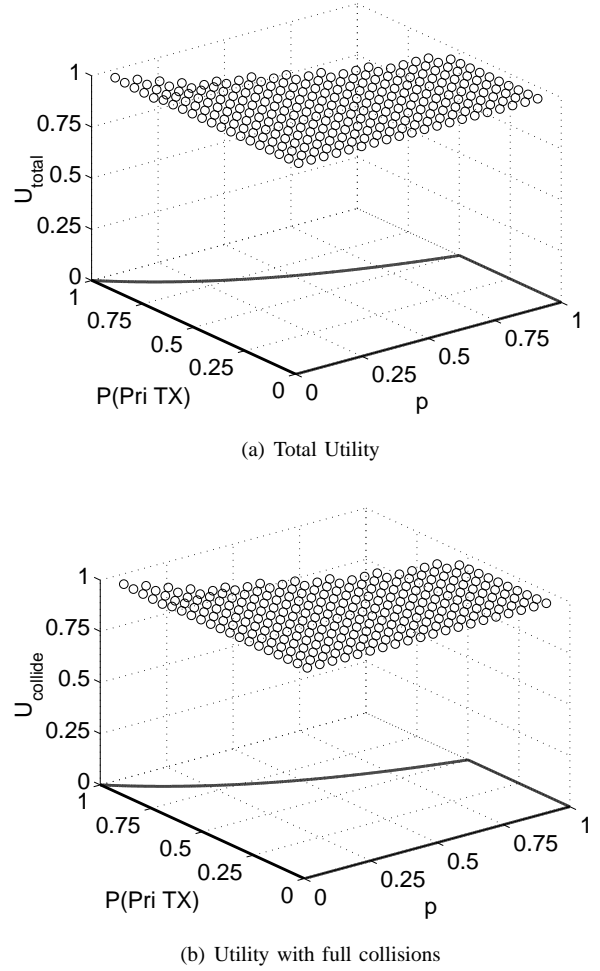


Fig. 16. With $p_{catch} = 0.8$, $p_{pen} = 0.6$, $\beta = 3$, and a better p_{FA} - p_{MD} curve, the overall utility is increased.

the penalty is greater than a threshold. This threshold was empirically found for a perfect detector to be the same as it was with the original detector, shown in Fig. 14. So to remove incentive to cheat, regardless of the detector, the penalty simply has to exceed this threshold.

However, it may be of interest to intentionally far exceed the threshold: better p_{FA} means better utility for the secondary, so by simply setting the penalty very high, we can encourage better secondary-sensing technology. This means that regulation does not have to set a particular standard p_{FA} and p_{MD} that all users must meet, and it is possible to simply incentivize better technology instead of having to certify every device and then recertify when the standard needs to be adjusted.

V. THE EFFECT OF WRONGFUL PUNISHMENT

Thus far we have considered only a single cognitive user and perfect detection by the primary. However, when many users or less than perfect interference detection are considered, identifying the particular cheating user may be difficult. So, in this section we want to get a sense of how the incentives must change when a user can be wrongfully punished.

Still assuming a single band, a user can be falsely accused only when the primary is transmitting and another user is




Fig. 17. With $p_{catch} = 0.8$, $p_{pen} = 0.6$, $\beta = 0$, a $p_{wrong} = 0.2$ (the probability of false accusation) causes the secondary to cheat more often than the original case in Fig. 3




Fig. 18. The relationship between β and p_{wrong} to assure there is no incentive to cheat for any channel usage characteristic. Here $p_{pen} = 0.4$ and different lines correspond to different values of p_{catch} . The circled points are empirically determined; the lines are drawn from an empirically determined equation. Note that as p_{wrong} approaches p_{catch} , the β required to assure no incentive to cheat goes to infinity.

cheating. We will assume for simplicity that there is always a malfunctioning unit. We can then model the probability of being falsely accused as a transition from S_2 (in Fig. 1) to the penalty boxes, controlled by a probability p_{wrong} . Intuitively, the secondary will be tempted to cheat more often if there is a chance of wrongful punishment because if it will be punished anyway, it may as well use the channel. Fig. 17 shows the same parameters as in Fig. 3 but with a 0.2 probability of going to the penalty box without actually cheating. Indeed, with the chance of wrongful accusation, the secondary will cheat in a larger set of channel usage scenarios.

To understand the incentives required to control cheating with wrongful punishment, we again find the bounding parameter values required to guarantee no incentive to cheat for any channel usage characteristic. These are plotted in Fig. 18. For illustration purposes, p_{pen} is fixed at 0.4, and the rest of the parameters are varied. Different lines correspond to different values of p_{catch} and show the β required for different p_{wrong} .

The circled points are empirically determined values, and the connecting lines correspond to

$$\beta = \frac{1 - p_{pen} + 0.9p_{wrong}}{p_{catch} - p_{wrong}} \quad (10)$$

which was empirically determined. Note that if $p_{wrong} = 0$, this reverts back to (8) found previously. Also, as p_{wrong} approaches p_{catch} , the β required to assure no incentive to cheat goes to infinity. If the regulator punishes everyone when interference is detected, as soon as one user begins to cheat no amount of punishment will incentivize the other users not to cheat as well.

VI. CONCLUDING REMARKS

This paper has introduced a very simple toy model for cognitive radio operation. Even with such a simple model, we can see some important effects:

- To incentivize cognitive users not to cheat, there must be a probability that they are caught and forced to undergo a penalty.
- It is not sufficient to penalize the cognitive user only by denying it access to the band that it is violating; some form of additional penalty is needed to deter cheating if we want a rule that works universally over primary usage characteristics.
- Banning the user from cognitive access to other bands as well can serve as such an additional penalty.
- The duration of the penalty must be set appropriately to incentivize proper behavior.
- Setting the penalty high enough gives rational cognitive users an incentive to develop appropriately sensitive detection algorithms. There is no need to regulate at the level of sensitivity itself.
- If the regulator punishes all cognitive users when one of them cheats, as soon as one cheats no level of punishment can incentivize the others not to cheat

This represents the start of an investigation. Much more should be explored, and the results of this model should be combined with a calculation of the overheads required to achieve the required level of enforcement.

REFERENCES

- [1] "Spectrum policy task force report," Tech. Rep. 02-135, Federal Communications Commission, Nov 2002.
- [2] R. W. Broderson, A. Wolisz, D. Cabric, S. M. Mishra, and D. Willkomm, "White paper: CORVUS: A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum," tech. rep., 2004.
- [3] M. A. McHenry and K. Steadman, "Spectrum occupancy measurements, location 1 of 6: Riverbend park, great falls, virginia," tech. rep., 2005.
- [4] G. Atia, A. Sahai, and S. Venkatesh, "Spectrum Enforcement and Liability Assignment in Cognitive Radio Systems," *Submitted to DySPAN*, 2008.
- [5] R. H. Coase, "The Federal Communications Commission," *Journal of Law and Economics*, vol. 2, pp. 1–40, October 1959.
- [6] R. Tandra, S. M. Mishra, and A. Sahai, "What is a spectrum hole and what does it take to recognize one?," *submitted to Proceedings of the IEEE*, 2008.
- [7] R. Etkin, A. Parekh, and D. Tse, "Spectrum sharing for unlicensed bands," *IEEE DySPAN*, Nov. 8-11 2005.

- [8] A. S. de Vany, R. D. Eckert, C. J. 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. 21, pp. 1499–1561, June 1969.
- [9] D. Hatfield and P. Weiser, "Toward Property Rights in Spectrum: The Difficult Policy Choices Ahead," *CATO Institute*, August 2006.
- [10] E. Goodman, "Spectrum Rights in the Telecosm to Come," *San Diego Law Review*, vol. 41, no. 269, 2004.
- [11] 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.
- [12] R. Tandra and A. Sahai, "SNR walls for signal detection," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, pp. 4–17, February 2008.
- [13] S. M. Mishra, A. Sahai, and R. W. Broderson, "Cooperative sensing among Cognitive Radios," *ICC*, June 6-10 2006.
- [14] S. M. Mishra, R. Tandra, and A. Sahai, "The case for Wideband Sensing," *Allerton Conference on Communication, Control, and Computing*, Sept. 26-28 2007.