

Crime and Punishment for Cognitive Radios

Kristen Ann Woyach and Anant Sahai
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley, CA
Emails: {kwoyach,sahai}@eecs.berkeley.edu

George Atia and Venkatesh Saligrama
Department of Electrical and Computer Engineering
Boston University, Boston, MA
Emails: {geokamal,srv}@bu.edu

Abstract—Frequency-agile radios hold the potential for improving spectrum utilization by allowing wireless systems to dynamically adapt their spectral footprint based on local conditions. Whether this is done using market mechanisms or opportunistic approaches, the gains result from shifting some responsibility for avoiding harmful interference from the static “regulatory layer” to layers that can adapt at runtime. However, this leaves open the major problem of how to enforce/incentivize compliance and what the structure of “light-handed” regulation should be. This paper addresses this question by developing a model for the incentives associated with cheating and for the tradeoffs between different elements of an enforcement structure which will effectively deter cheating. It then investigates a code-based scheme for detecting and assigning liability to culprits.

I. INTRODUCTION

The Federal Communications Commission currently sits at a crossroads, deciding what regulation is going to look like for the next generation of wireless devices [1]. The current command-and-control model, in which spectrum is parceled and allocated to specific uses and companies, was designed for broadcast systems such as TV and AM/FM radio. However, as technology changes, this approach becomes less applicable. The current centralized solution has difficulty managing allocations on the heterogeneous usage scales of interest and so leaves “holes” in both time and space where valuable spectrum is being wasted [2]. Indeed, although spectrum looks scarce to anyone who wants an allocation [3], anyone actually taking measurements realizes that most of the available spectrum is in fact underused [4].

Scholars have debated how to solve this problem. While all agree that decentralized and more “light-handed” regulation is desirable, the form of this regulation is contested. Spectrum privatization advocates rely on market forces to determine not only who will be allowed to transmit but also the size of the parcels of spectrum they are allotted [5], [6]. In this model, government regulation certifies devices, monitors market transactions, and resolves disputes as civil offenses through the courts. Spectrum commons advocates note that with current technological advances, a simpler approach is possible that puts the burden of regulation entirely on equipment — any certified device may transmit [7].

Regardless of the regulatory intention, the introduction of frequency-agile and software-defined radios brings opportunistic use into the picture. Cognitive radios are autonomous and possibly adaptive, allowing them to adjust their transmis-

sion patterns according to local observations [8]. So, even if a company officially has primary rights to a piece of spectrum, these agile radios will be capable of detecting unused portions and using it for short periods of time. Whether they will do so legally is yet to be seen, but that they will try is inevitable.

The decision to support opportunistic use is therefore really a decision about the circumstances under which a primary user of a band is allowed to issue a cease-and-desist order. If opportunistic use is to be discouraged, the primary user would be allowed to issue such an order whenever it detects any use that is not its own. However, this encourages spectrum owners to behave as “spectrum trolls” [9]. Dubbed by Hatfield, these users do not productively use their allotted spectrum but rather wait for others to use it unlawfully and then threaten legal action if the opportunistic users do not pay a bribe. On the other hand, if opportunistic use is to be supported, primary users must demonstrate harmful interference before issuing a cease-and-desist order.

In this paper, we take a stance encouraging opportunistic use because not only does it allow better utilization of spectrum, but legalizing opportunistic use also permits rational regulation and certification of the radios used for this purpose.

However, certification is not simple. For increasingly autonomous and mobile radios, frequency-agility runs the risk of being the wireless equivalent of Plato’s Ring of Gyges. Faulhaber raises this specter through his discussion of “hit and run radios” that are virtually uncatchable because they turn on, use the spectrum for a period of time, and turn off without a trace [10]. We cannot ignore the important problem of whether/how sharing rules for dynamic spectrum-access schemes can be enforced.

The current literature indicates that by using game theory, one can show that equal users can self-enforce [11]–[13] by balancing their own quality of service against the interference they are causing other users, allowing for a range of stable and fair equilibria. However, this self-enforcement breaks down when users are not equal. Consider a case with two users; the first can cause very little interference to the second while the second can cause a great deal of interference to the first. The first has neither defense nor ammunition. Without a possibly external force to which the second is vulnerable, the first cannot reasonably believe that the second will follow sharing rules. Indeed, vulnerability is the mother of trust.

In much of the spectrum commons and spectrum-sharing literature, regulation to help support unequal users appears

as a purely *a priori* device-level certification [7], [14], [15]. Either the hardware, the software, or both, are certified to meet certain standards that facilitate coexistence before they can be deployed. For a radio using only its own observations, lab tests are sufficient to certify that it is properly following sharing rules. But evidence in [2] and [16] suggests that a solo radio’s detection capabilities are not effective, and so cooperative detection schemes are desirable to enable coexistence. Unfortunately, certifying that a cooperative network is correctly following sharing rules seems difficult and may require digging through thousands of lines of code to make sure an adversarial vendor is not trying to fool the test.

Even with proper certification, some run-time policing is still necessary. Devices can misbehave both by maliciously trying to cause harm and by inadvertently malfunctioning. The wireless medium is such that there is no natural protection against malfunctioning nodes. Indeed, these devices can cause a great deal of harm and must be externally stopped in a relatively short amount of time to maintain quality of service for all other users. So, certification and device recalls are insufficient; policing and a properly certified kill-switch are required to stop malfunctioning nodes at runtime [15].

Given that some runtime enforcement is likely to be necessary, can it be used to make certification simpler? A kill-switch is effective but extreme; it is the radio analog of a death sentence. Therefore, it can only be applied when guilt is undeniable (as with a malfunctioning node that is always transmitting despite sharing rules). It also suggests that we should be thinking of spectrum offenses like criminal instead of civil offenses and therefore should support an analogous gradation of punishment. If a kill-switch is already necessary, it is not much harder to also certify a temporary time-out or jail sentence. A jail sentence presents a level of punishment that can be applied repeatedly, even in slightly ambiguous cases. The vulnerability to this credible threat can be used to deter intentional misbehavior.

In this paper, we consider a jail-based enforcement scheme to explore the incentives involved with cheating and what punishment is necessary to deter it. We then explore the overhead required for a catching/punishing scheme and when it is in the cognitive user’s best interest to opportunistically use unoccupied primary bands. Finally, we consider what is the minimum certification required to allow policing to be effective. To this end, we introduce a scheme to give radios identities that facilitate the catching and punishing of misbehaving nodes.

II. SINGLE BAND

In developing this game-theoretic model, we make several simplifying assumptions: we first assume that the spectrum “holes” exist in time, so we model the primary usage in each band as a two-state Markov chain with probabilities p and q of turning off and on, respectively. We also assume that the transmission characteristics of both primary and cognitive user are slotted and synchronized, ignoring sensing time. If the cognitive user is transmitting at the same time as the primary, it is considered to be cheating and so has a chance

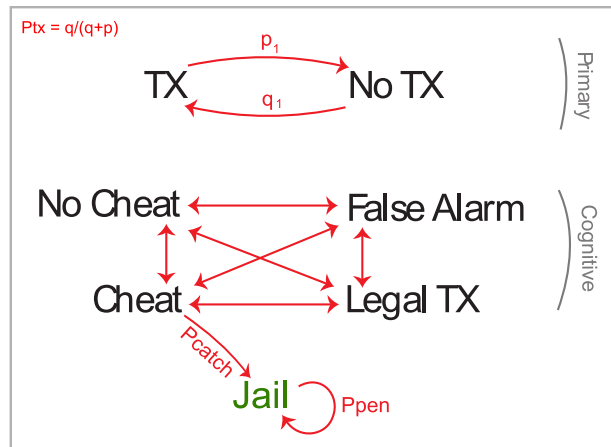


Fig. 1. Markov model for enforcement in a single band: the primary user follows a transmission pattern characterized by the two-state chain at the top. In response to this usage, the cognitive user chooses between the actions in the bottom chain. When the cognitive user is cheating (i.e. using the band when the primary is also transmitting), it has a probability of being sent to jail. While in jail, the cognitive user is unable to use the band until a timer (determined by P_{pen}) runs out.

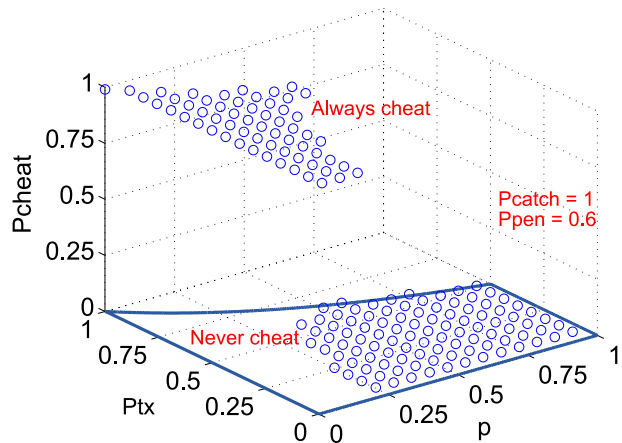


Fig. 2. Typical cheating behavior for a secondary user. If it is worthwhile to cheat, it is worthwhile to always cheat. The choice of whether to cheat is a function of the probability of being caught, the duration of jail sentences, and the transmission characteristics of the primary. If the primary is always transmitting, jail as a time-out is not an effective deterrent against cheating because an honest cognitive user would not get to transmit anyway.

of being punished. These assumptions are made to simplify the analysis while still retaining the important interactions between the primary and cognitive users.

We model punishment as a “spectrum jail” system in which the cognitive user loses the privilege to use the band if it is caught cheating. The interaction between primary and cognitive users for a single band, then, can be modeled with the Markov chain depicted in Fig. 1. The primary usage is characterized by the two state chain at the top, which instigates the reactions of the cognitive usage chain at the bottom. The primary user’s on-off status determines movement horizontally through the chain. When the primary user is active, the secondary may either be legally transmitting, or seeing a False Alarm, with corresponding probabilities of $1 - P_{FA}$ and P_{FA} . When the primary user is transmitting, the




Fig. 3. These lines describe how painful jail must be in order to deter cheating. If β is above these line corresponding to the other enforcement parameters, there is no incentive to cheat, regardless of the primary transmission pattern.

cognitive user can choose to cheat with probability P_{cheat} . If it is in the cheating state, the cognitive user will go to jail on the next step with probability P_{catch} ¹ (the probability of being caught). Once in jail, the secondary must wait there for an amount of time determined by P_{pen} before it is allowed to rejoin the game.²

In this game, the cognitive user adjusts its probability of cheating to maximize its utility, defined as the average amount of time it is transmitting.³ So, its objective is:

$$\max_{P_{cheat}} U = \max_{P_{cheat}} \pi_{legal} + \pi_{cheat}, \quad (1)$$

where U is the utility gained; π_{legal} and π_{cheat} are the stationary probabilities of legally transmitting and cheating, respectively. The primary (or its proxy the regulator) on the other hand, wants to minimize the time the secondary is cheating using its adjustable parameters, P_{catch} and P_{pen} . Note that the regulator must use these parameters in slightly different ways: P_{catch} is constrained by the catchability of cheating cognitive users and the deployment density of monitors. As conditions change during run-time, this parameter can be adjusted. The effectiveness of P_{pen} , however, relies on cognitive users respecting the jail-time. Therefore, it must

¹ P_{catch} captures two distinct effects. The first is the primary user's imperfect catching mechanism, which will miss some cheating cognitive users. It also captures the case when the signal between primary and cognitive users is faded. So even though the cheater is caught, it may not hear the "go to jail" command and thus it could continue to transmit.

²In the real world, we generally think of jail sentences as being for a specific, deterministic duration. Here, time in jail is considered as a probabilistic quantity with the correct mean to simplify the analysis.

³We assume that the cognitive user gets the utility of using a clean band when cheating despite interference from the primary user's transmission. This is done to account for the case that is most worrisome to a primary user: when the path from the primary transmitter to the cognitive user is badly shadowed, but the path from cognitive user to the primary receiver is not. So, the primary user can cause little damage to the cognitive user, but the cognitive user can cause major damage to the primary user.

be set beforehand with a P_{catch} in mind, and devices must be certified to respect particular jail sentences.

Fig. 2 shows the typical behavior of the secondary user, over different primary transmit characteristics when $P_{catch} = 1$ and $P_{pen} = 0.6$. When the primary is rarely active and switches quickly from transmitting to not transmitting (low $P_{TX} = q/(q+p)$ and high p), the secondary has no incentive to cheat, and so $P_{cheat} = 0$. However, if the primary user is nearly always active or switches slowly, the cognitive user will be more tempted to cheat. In fact, if the primary user is always transmitting, the secondary user will always be tempted regardless of the values of P_{pen} and P_{catch} because it costs the same amount to sit in jail as it does to wait for the primary user to turn off.

Because the regulator must set the requisite P_{pen} at certification time, the chosen parameters must work for any p and q . So, the regulator sets them with respect to the worst case: when the primary user is always transmitting. As noted above, P_{catch} and P_{pen} are insufficient to deter cheating in this case because the secondary user gains utility by simply bouncing in and out of jail. Therefore, we introduce a factor β which is the cost of sitting in jail. In the real world, β would correspond to an extra punishment, such as a fine or the emotional/physical hardship of being in jail. The secondary objective is now

$$\max_{P_{cheat}} U = \max_{P_{cheat}} \pi_{legal} + \pi_{cheat} - \beta \pi_{jail}. \quad (2)$$

Using the extra β factor, the primary can set P_{pen} and P_{cheat} to account for the worst case of the primary always transmitting by considering a simple two-state Markov chain. The secondary goes to jail with probability $P_{cheat}P_{catch}$ and leaves with probability $1 - P_{pen}$. The regulator should then set the parameters such that

$$P_{cheat}\pi_{not_jail} - \beta\pi_{jail} < 0, \quad (3)$$

or the utility lost by being in jail is greater than that gained by cheating for any value of P_{cheat} . This leads to a condition

$$\beta > \frac{1 - P_{pen}}{P_{catch}} \quad (4)$$

to dissuade cheating regardless of transmission pattern. The boundary for this condition is shown in Fig. 3.

However, an extra punishment for sitting in jail may be problematic. Although it is relatively easy to implement and certify a "shut down" command, fines would require extra overhead in terms of a government-certified billing system. They also have other, less obvious effects such as enabling Hatfield's "spectrum trolls" [9]. An alternative is to have β payable in kind so that it requires no additional infrastructure and does not encourage destructive behavior. Consequently, we propose considering cognitive radio as a bandwidth expander in which each user has a dedicated home band of value β and may expand into other bands by staking its home band against unlawful use. This interpretation lends itself naturally to expansion into several bands. Therefore, we now consider the multiband case.

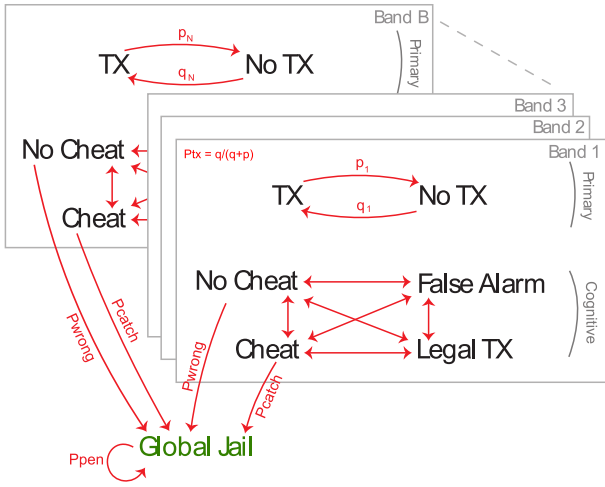


Fig. 4. Depiction of the Markov chain for multiple bands. Here, each band has an independent primary user with an independent transmission pattern. The cognitive user may choose how to cheat on each band independently, but if it is caught cheating in any of the bands, it is sent to a Global Jail. While in jail, the cognitive user cannot transmit in its home band or in any of the other bands.




Fig. 5. Requirement on the number of home bands to deter cheating regardless of primary transmission pattern when multiple bands are considered. Each new band introduces one more possible unit of utility, and thus more temptation. The required β is now B times that required for a single band.

III. MULTIPLE BANDS WITH GLOBAL JAIL

A. Perfect Justice

The multiple-band case can be modeled with the Markov chain depicted in Fig. 4. Each band has an independent primary with separate characteristics. The secondary can choose whether to cheat separately in each band, but if it is caught cheating in even one band, it goes to a Global Jail, where it is not allowed to use either the expansion bands or its home band. Note that the Global Jail is a necessary consideration with multiple bands because if each band had a separate jail the multiple-band case would devolve into several single-band problems.

For now, assume that the regulator has no uncertainty and so will catch the correct offending party with probability




Fig. 6. The required β to deter cheating rises dramatically when $P_{wrong} \neq 0$. Of particular note is the behavior when P_{wrong} is very close to P_{catch} . If you will be sent to jail with the same probability anyway, you might as well cheat and get some utility for it.

1. P_{catch} is now the probability that the cognitive user hears the “go to jail” command despite a faded wireless channel. Perfect detection of cheaters also implies that the probability of being wrongfully convicted (denoted P_{wrong} in the Markov chain) is zero. To simplify the analysis, we assume that the bands are identical: the primary transition probabilities p and q , although independent, are the same for all bands.

In this case, we can derive a bound on the necessary number of home bands required to dissuade cheating by considering the worst case of all the primaries always transmitting. Solving for the β required for the cost of jail to be greater than the utility gained by cheating, we find the condition

$$\beta > B \frac{1 - P_{pen}}{P_{catch}} \quad (5)$$

where B is the number of bands the secondary is capable of expanding into. The boundary of this function is shown in Fig. 5. This equation is intuitively pleasing as the temptation to cheat should scale with the number of opportunities to cheat. This can alternatively be thought of as a condition on P_{pen} :

$$P_{pen} > 1 - \frac{\beta}{B} P_{catch} \quad (6)$$

because the number of home bands is presumably fixed before the enforcement parameters are set.

Note that P_{pen} is getting closer to 1 as the number of expansion bands grows. However, the honest cognitive user will never be sent to jail in this model, so the rising P_{pen} does not hamper further expansion in any way. In reality, there is a chance that innocent cognitive users are sent jail. We now consider the effect of this uncertainty.

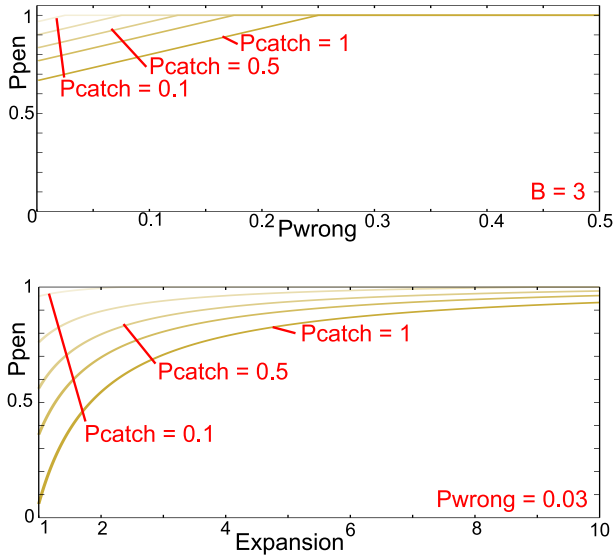


Fig. 7. The regulator sets P_{pen} such that for a given value of home band, β and a particular expansion B




Fig. 8. As P_{pen} grows, so does the amount of time spent in jail for wrongful convictions. This influences the utility an honest cognitive user will gain from extra bands. Shown here is the cognitive user's utility as a function of the number of expansion bands, along with the percentage of time the user spends in jail. The utility peaks around the time the user spends 50% of its time in jail.

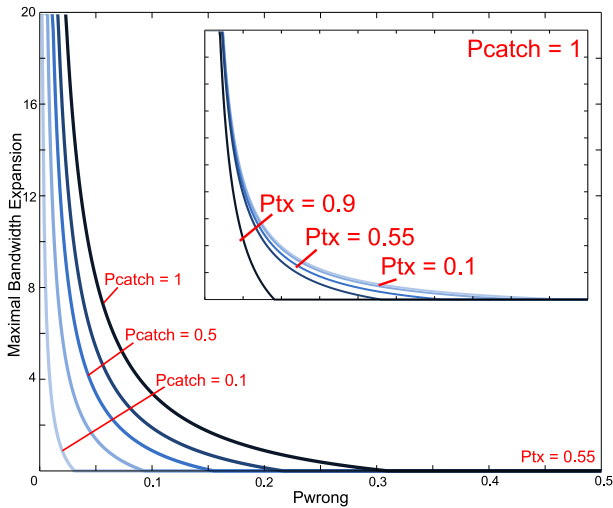


Fig. 9. Assuming P_{pen} is as low as possible to satisfy the primary, the secondary can choose how many bands to expand into to maximize its overall utility. Notice that for expansion to be large, P_{wrong} must be very small.




Fig. 10. We define overhead as the proportion of utility available that the cognitive user cannot gain due to spending time in jail. When plotted against the B which maximizes utility, we see that the cognitive user will never have an overhead greater than 50%.

B. Imperfect Justice

We model regulator uncertainty with the parameter P_{wrong} in Fig. 4. This parameter captures two effects: if the cognitive user misses detecting the primary and uses the band, it will get sent to jail even though it is not intentionally cheating. In this case, the effective P_{wrong} could scale to 1 as the number of bands increases. P_{wrong} also captures uncertainty with the regulator if it mis-identifies the cheating user, or collectively punishes many cognitive users for one misbehaving node. In either of these cases, the secondary will be sent to jail innocently with some probability that does not scale with the number of bands.

Here we consider just such a P_{wrong} that does not scale with the number of expansion bands. As before, the regulator wants to set the enforcement parameters so that there is no incentive to cheat even if the primary users are all transmitting all the time. Again, we analyze this as a two-state Markov chain with the probability of going to jail $P_{cheat}P_{catch} + (1 - P_{cheat})P_{wrong}$ and the probability of leaving jail $1 - P_{pen}$. Then, the primary should set the parameters so that

$$BP_{cheat}\pi_{not_jail} - \beta\pi_{jail} < 0, \quad (7)$$

or the utility gained by cheating is less than the utility gained by not. This produces the condition:

$$\beta > B \frac{1 - P_{pen} + P_{wrong}}{P_{catch} - P_{wrong}}. \quad (8)$$

The boundary of this condition is shown in Fig. 6 for $B = 3$, $P_{pen} = 0.6$, and different values of P_{catch} . The interesting thing to note here is that as P_{wrong} approaches P_{catch} , the necessary β goes to infinity. No amount of extra punishment can deter cheating if you will be sent to jail with the same probability anyway.

As before, we should be considering this condition as the P_{pen} required for a given number of home and desired number of expansion bands:

$$P_{pen} > 1 + P_{wrong} - \frac{\beta}{B}(P_{catch} - P_{wrong}). \quad (9)$$

P_{pen} is again rising with B . Now, however, the cognitive user cares how high P_{pen} becomes because it will be sent to jail with some probability despite honest use.

Consider the following: for a given number of home bands, β , and a particular P_{catch} , the primary sets P_{pen} in (9) as low as possible as a function of the bandwidth expansion B and P_{wrong} . The result is shown in Fig. 7. As P_{wrong} grows, P_{pen} goes to 1.

Likewise, for a given P_{wrong} , as the secondary chooses to expand into more bands, the P_{pen} required to assuage the regulator goes up. Spending more time in jail affects the secondary user's utility as shown in Fig. 8. The utility peaks around the point where the cognitive user spends half its time in jail, and then falls, becoming negative at some point.

This utility function determines the game that the cognitive user is willing to play. If the cognitive user can choose only between not playing or playing with a particular number of expansion bands, the zero-crossing of the utility function is important. The cognitive user will play if the number of bands results in positive extra utility, and not otherwise. If the cognitive user can choose how many bands it will expand into, it will choose the number of bands that maximizes its utility. So, P_{wrong} dictates that even a greedy secondary user will choose to expand only so far into other bands.

To make this utility maximization more precise, consider the secondary user operating on a two-state Markov chain with states Jail and Not Jail and a probability of going to jail P_{jail} . The probability of leaving jail is $1 - P_{pen}$. The secondary wants to solve the following optimization problem:

$$\max_B \frac{1 - P_{pen}}{P_{jail} + 1 - P_{pen}} U_{noTX} - \beta \frac{P_{jail}}{P_{jail} + 1 - P_{pen}} \quad (10)$$

where P_{pen} is set to have equality in (9), and U_{noTX} is the utility gained when the primary is not transmitting.

Assume that at run-time, the secondary user estimates the average probability that the primary users are transmitting ($P_{TX} = q/(q+p)$). It is then allowed to adjust its number of expansion bands, and corresponding P_{pen} , to maximize its utility. However, the time spent in jail does not depend just on the average probability of primary transmission; it depends on how quickly the primary users transition between transmitting and not. If the primary users switch very quickly, the cognitive user will spend more time in jail relative to slowly switching primary users. So, the worst case is when the primary users bounce back and forth at every time step. In maximizing its utility, the cognitive user should account for the worst case. However in real situations, primary users are more likely to stick in particular states rather than simply bouncing back and forth. Therefore, it is sufficient to calculate the maximum utility for the case when the primary usage is iid at each time step ($p+q=1$).

With the iid primary usage assumption, $U_{noTX} = B(1 - P_{TX})$ and $P_{jail} = P_{wrong}(1 - (1 - P_{TX})^B)$. ($1 - (1 - P_{TX})^B \approx 1$ for most cases of interest, so we will approximate P_{jail} as P_{wrong}). The optimal bandwidth expansion B/β is then

$$\frac{B}{\beta} = \frac{P_{catch}(1 - P_{TX}) + P_{wrong}(P_{TX} - 2)}{2P_{wrong}(1 - P_{TX})}. \quad (11)$$

This function is plotted in Fig. 9, varying P_{catch} in the larger plot and varying P_{TX} in the cut-out. Note that for large bandwidth expansions, P_{wrong} must be very small, and small changes produce large variations in the maximal expansion. Therefore, the secondary has an incentive to create the best detector possible to keep P_{wrong} as low as possible.

We would also like to get a sense of the overhead incurred by this punishment scheme. We are taking the stance of encouraging opportunistic use to fill the "spectrum holes" and so the overhead is the percent of available bandwidth that the cognitive user is not able to use because of jail time. Therefore the overhead is defined as

$$Overhead = \frac{\beta + BP_{TX} - (U + \beta)}{\beta + BP_{TX}}, \quad (12)$$

where U is the utility being maximized in (10). The maximal expansion vs. overhead is plotted in Fig. 10. This figure gives a guide for the expansion possible for a given amount of allowed overhead and a particular P_{wrong} . Notice that the amount of overhead never exceeds 0.5 for desirable amounts of expansion.

IV. IDENTIFYING SPECTRUM VIOLATORS

So far, the only certification requirement we have imposed is that radios obey a wireless "go directly to jail" command and stay there for an expected duration of $\frac{1}{1 - P_{pen}}$. This time depends on the number of bands B that the frequency-agile radio can expand into as well as a lower bound on P_{catch} and an upper bound on P_{wrong} . To reduce the overhead that they must pay, rational radio designers *want* to make their radios catchable and demonstrate as much to the regulators at device certification time.

However, the plots in Figure 9 and the expression in (11) reveal that it is P_{wrong} that has the most significant effect because it occurs in the denominator. This is the cognitive radio analog of the popular sentiment "better a hundred guilty men go free than a single innocent man go to jail" that motivates our "innocent until proven guilty beyond reasonable doubt" criminal justice system. To reduce the P_{wrong} , it is important to both reduce the radio system's fear of inadvertently causing harmful interference [2] as well as to make sure that the radio system has its own identity distinct from those of other systems. This is so that it does not incur wrongful convictions due to being caught in a dragnet along with other devices that are causing interference or due to mistaken identity.

One aspect of a distinct identity is for the radio to have a way to reliably receive an individualized message telling

it to go to jail. For this aspect, it is clear that traditional communications thinking applies: the outage-capacity (for a suitably low probability of outage to keep P_{catch} large) of the regulatory control channel, together with the responsiveness delay, will determine the number of distinct identities that can be supported.

The other aspect of identity is that which enables a primary user to point its finger at the relevant culprit or culprits. There are many potential approaches to ‘identity.’ In the most straight-forward approach, identity is explicitly transmitted by the physical layer as a separate wireless signal in a mandated format. If a primary user experiences harmful interference, then it merely has to decode this signal to learn the identities of all the potential interferers so that they can be penalized. Such an “identification beacon” would necessarily impose an overhead on the secondary users and one could analyze the possible tradeoffs. However, while this approach is conceptually simple, it has four major shortcomings:

- 1) It forces us to mandate a standard PHY-layer format for *transmission* of this identity information. This adds additional complexity⁴ to systems that want to use a different format/modulation for their own signals. Moreover, this format would have to be tamper-proof and thus might impose costly re-certification requirements for changes that could otherwise be handled as simple software updates.
- 2) It imposes a decoder PHY burden on the primary user to implement a way to decode this identity information so that it can penalize secondary systems that are cheating in the vicinity. This is in addition to the primary’s own PHY layer for decoding its own data. Sadly, if the regulation is successful and the threat of punishment is enough to prevent cheating, then this particular part of a primary system will be unexercised and hence is likely to suffer “bit-rot” as the primary systems evolve.
- 3) It does not allow the primary user to distinguish between harmful interference and unfortunate fading or bad luck. A primary user might simply be out of range of its transmitter or it might be drowning in harmful interference. There is no way to tell them apart if the secondary identity information was carried just by a separate beacon.
- 4) A broadcast identity does not distinguish between the guilty and the innocent bystanders. Thus it greatly reduces the incentive to deploy innovative approaches to reduce interference. For example, a cognitive-radio network might be able to use beamforming to null out its transmissions at the primary receiver. However, if any other cognitive radio causes harmful interference, the careful radios will also be punished since they were also in the neighborhood at the time of the incident.

A second approach to identity is developed in [17]–[19]

⁴Since the transmit and receive front ends are generally distinct in radio systems, the fact that we already have to support the decoding of a standard, regulatory control channel PHY-layer on the receive side does not mitigate this extra burden.




Fig. 11. Taboo-based identities, demonstrated here as the composition of three levels: network, user, and device. The taboo times are different in different bands to enable intelligent frequency hopping to maintain steady low-latency links.

where idiosyncrasies of the radio front-ends are used to identify individual devices. This avoids the need for an explicit beacon, but also requires good interference signal strength because these idiosyncrasies were never designed to differentiate different radios in signal space. In addition, while this “accidental identity” approach deals with the first objection above, the other three remain.

Furthermore, such accidental identities provide no way of having multiple coordinates associated with a single transmission. For example, a transmission might originate from a particular device that is a part of a particular network and being used by a particular human user. An explicit beacon could just concatenate bit-fields to transmit all three identities but there is no way to do this with accidental identities. For example, contrast “tall female, short blond hair, slim build, wearing a purple bodysuit” as a description of a humanoid suspect with a more engineered identity like “Seven of Nine, Tertiary Adjunct to Unimatrix Zero-One.”

Figure 11 shows another approach that is the wireless analog of “Geographic Profiling” of criminals. It has been observed that serial killers tend to maintain a taboo buffer zone around their homes wherein they do not kill anyone [20]. The wireless equivalent of this taboo-based identity is to design an identity-specific code that specifies which time-slots are taboo for this “temporal profile” identity. This taboo can easily be certified in the hardware and different identities can be stacked by giving each code a veto over using slots that are taboo to it. With this approach, the identity of a device is implicitly announced by the pattern of interference itself! This avoids all the problems above: no separate PHY is needed, there is no additional decoder burden on the primary since all it needs to record is the pattern of interference, and only the secondary users that are actually causing interference will show up.

The cost of this taboo-code approach is that certain slots are not available for use. A preliminary analysis of the required slot overhead is given in [21], [22].

V. CONCLUSIONS

Once radios are “cognitive” and have a degree of autonomy, it is unrealistic to rely purely on *a priori* certification to guarantee safe behavior. Instead, it makes sense to certify that the radios are appropriately vulnerable to a degree of enforcement that can take place at runtime. The vulnerability should be such that rational radios have no incentive to

cheat and can thus be trusted. In this paper, we explore the imposition of a two-part obligation:

- Radios should obey a “go directly to jail command” in which they are blocked from any wireless transmissions for a certain sentence that depends on the current capability of the radio for bandwidth expansion as well as the value of the home band they are able to stake.
- Radios should have their transmissions obey a “temporal profile” in which each radio has an individualized sequence of band-specific time-slots that are taboo to it.

These two obligations seem to be easy to certify and suffice to guarantee that no rational radio will intentionally cheat.

This paper is the first in a trilogy of conference papers to introduce these concepts. This one focuses on the threatened punishment required to deter cheating as well as the performance required from the identity system. The identity system itself is discussed in some detail in [21] with a focus on an idealized deterministic model of interference. The model for interference is made probabilistic in [22]. Even so, these three papers represent just the beginning of a research effort. Much more needs to be done before this idea is implementable in practice. For example, practical tractable code families are required as well as concrete approaches to certification tests.

ACKNOWLEDGMENTS

This invited paper incorporates, in Section IV, a few passages directly from [21] the work on which was supported in part by PECASE grant no. N00014-02-100362, NSF CAREER award ECS-0449194, the MIT-Portugal program, Sumitomo Electric, and NSF Grants CCF-0430983, CNS-0435353, ANI-0326503, CCF-0635372, and CNS-0627161. The new results here relating to incentives and punishment were supported by an NSF Graduate Research Fellowship. The presentation here was shaped by the much higher-level presentation in [16] and is thus indebted to discussions with Rahul Tandra and Mubaraq Mishra.

REFERENCES

- [1] “Spectrum policy task force report,” Tech. Rep. 02-135, Federal Communications Commission, Nov. 2002.
- [2] 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.
- [3] NTIA, “U.S. Frequency Allocations.” Available online: <http://www.ntia.doc.gov/osmhome/allochrt.pdf>.
- [4] M. A. McHenry and K. Steadman, “Spectrum occupancy measurements, location 1 of 6: Riverbend park, Great Falls, Virginia,” tech. rep., 2005.
- [5] R. H. Coase, “The Federal Communications Commission,” *Journal of Law and Economics*, vol. 2, pp. 1–40, Oct. 1959.
- [6] 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.
- [7] 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.
- [8] J. Mitola, *Cognitive Radio: an integrated agent architecture for software defined radio*. Ph.d. thesis, KTH Royal Inst. of Tech., Stockholm, Sweden, 2000.
- [9] D. Hatfield and P. Weiser, “Toward Property Rights in Spectrum: The Difficult Policy Choices Ahead,” *CATO Institute*, Aug. 2006.

- [10] G. R. Faulhaber, “Wireless telecommunications: Spectrum as a critical resource,” *Southern California Law Review*, vol. 79, Mar. 2006.
- [11] 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.
- [12] C. Rose, S. Ulukus, and R. D. Yates, “Wireless systems and interference avoidance,” *IEEE Transactions on Wireless Communications*, vol. 1, pp. 415 — 428, July 2002.
- [13] D. Popescu, O. Popescu, and C. Rose, “Interference avoidance versus iterative water filling in multiaccess vector channels,” *IEEE 60th Vehicular Technology Conference*, Sept. 2004.
- [14] 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.
- [15] W. Xu, P. Kamat, and W. Trappe, “TRIESTE: A trusted radio infrastructure for enforcing spectrum etiquettes,” *First IEEE Workshop on Networking Technologies for Software Defined Radio Networks*, Sept. 2006.
- [16] A. Sahai, S. M. Mishra, R. Tandra, and K. A. Woyach, “Cognitive radios for spectrum sharing,” *IEEE Signal Processing Magazine*, Jan. 2009.
- [17] V. Brik, S. Banerjee, M. Gruteser, and S. Oh, “PARADIS: Physical 802.11 device identification with radiometric signatures,” in *Proceedings of ACM Mobicom*, (Burlingame, CA), Sept. 2008.
- [18] J. Hall, M. Barbeau, and E. Kranakis, “Radio frequency fingerprinting for intrusion detection in wireless networks,” in *Defendable and Secure Computing*, 2005.
- [19] K. B. Rasmussen and S. Capkun, “Implications of radio fingerprinting on the security of sensor networks,” in *Proceedings of IEEE SecureComm*, 2007.
- [20] D. K. Rossmo, *Geographic profiling: target patterns of serial murderers*. Ph.d. thesis, Simon Fraser University, 1995.
- [21] 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.
- [22] G. Atia, V. Saligrama, and A. Sahai, “Spectrum enforcement and liability assignment in cognitive radio systems,” in *Proceedings of the 42nd Asilomar Conference on Signals, Systems and Computers*, Nov. 2008.