

Why the caged cognitive radio sings

Kristen Woyach and Anant Sahai
Department of EECS, UC Berkeley
Email: {kwoyach, sahai}@eecs.berkeley.edu

Abstract—In our earlier papers [1] and [2], we have proposed a jail-based enforcement mechanism for cognitive radios inspired by the human criminal justice system. In the previous papers, we covered a throughput-greedy cognitive user, and in this paper, we extend those results to devices that care about energy, as well as mixed devices that care about both energy and time. We do this by introducing a ‘singing’ sanction that forces devices to burn energy while they sit in jail.

Through this exploration, we see a number of effects coming out: a homeband (which may be an unlicensed band) is required to present an alternative band to legally transmit in when it is difficult to operate legally in the cognitive band. Also, when the primary is very rarely active, it is practically impossible to deter cheating, so alternate policy decisions must be made for these cases. Finally, it is possible to create a singing plus jail sanction that is sufficient to deter bad behavior for all types of devices in all types of situations. But in order to enforce against everyone, while keeping overhead low, the rate of wrongful convictions must be kept small.

I. INTRODUCTION

Our goal in [1], [2], and this paper is to understand what is required to build an enforcement system for cognitive radios that:

- 1) Is easy to certify
- 2) Is applicable to all current and future technology
- 3) Causes rational cognitive devices to behave like an honest user that does not cause harmful interference
- 4) Does not overly burden radios that would follow sharing rules anyway.

To do this, we are designing an enforcement system inspired by the criminal justice system that relies on run-time policing. In these papers, we are identifying corner cases and designing the system for those assuming that other cases will naturally be covered as well. In [1] and [2], we covered the corner case of radios that were throughput-greedy. However, there exist radios that do not care at all about throughput; they care about energy. In this paper, we extend our enforcement model to account for these energy users.

A. Background

In September 2010, the FCC released revised requirements for white space devices operating in the TV bands [3]. However, it did not specify exactly how the rules would be enforced and devices certified.

One possible implementation for certification is to test in a laboratory that a single radio, operating independently, is capable and willing to either check a database at the required intervals or sense a primary signal received at a level of -114dBm. Check, then, that this functionality cannot be

changed after deployment, and the radio will not be physically capable of violating the rules.

Other ideas for cognitive radio enforcement follow a similar pattern. For example, [4] rely on a certifiable etiquette module within each cognitive device to control its sensing and transmission. [5] uses a trusted token-system to allow radios to request/pay for the right to transmit at any given time.

But what about different bands or future technology? In a cellular band, where primaries are mobile, would secondaries be required to consult a database of real-time holes before every transmission? For sensing solutions, [6] has shown that sensing is more effective when networks of devices cooperate. These gains cannot be realized under a simple single-radio sensitivity rule. One could even imagine a secondary that would help relay for the primary, [7]. But how could we properly prove that it does what it should?

In [1] and [2], we have proposed a light-handed solution to the certification and enforcement problem based on run-time policing. We will recap and expand the arguments for this idea here. For more details and references, see [1] and [2].

Any regulation problem boils down to how to optimally combine the following components:

- 1) Certification requirements
- 2) Standards bodies and network effects
- 3) Run-time policing
- 4) Legal system and the courts

The simplest white-space enforcement achieves its objectives using certification alone, ignoring the other regulation components. If we knew the best possible technical solution to spectrum sharing, forcing all radios to use this technology would make sense. Unfortunately, we do not. Worse, past history says it is difficult to change the rules [8], creating an almost insurmountable barrier to entry for new technologies.

It has also been suggested that defining loose rules, and using the courts to correct them later [9] is a possibility. Given the difficulty of addressing interference under current common law [10], and given that the rules must be able to apply to both legacy and much more advanced technology at the same time, this is not an obviously good solution.

Is there another option worth exploring?

We must first ask what we want from regulation. We propose that it should be light-handed, which means (a) universal across all current and future technologies; (b) does not unreasonably impede the operation of such technologies; (c) is flexible to possible changes in sharing rules; and (d) is simple to certify, thereby lowering the barrier to entry.

A good model for such a system already exists: the criminal justice system. Consider, for example, the laws surrounding driving. Certification involves testing the car and issuing

drivers' licenses. Standards exist that most people follow such as passing on the left. Traffic cops monitor for unsafe drivers breaking the law. Courts are used to arbitrate disputes or prosecute particularly bad offenses. Finally, any changes in the existing laws governing the road require education of the people involved but no architectural changes to the system.

We are interested in understanding what a similar system for cognitive radio enforcement would look like and whether it could work. To port the idea from the human-realm to the radio-realm, we need to understand the necessary components and how radios differ from people.

There are two parts to a policing-based enforcement scheme: catching the criminals and punishing them appropriately. We will discuss briefly what we believe necessary to catch an offending radio, paying particular attention to how well the system catches actual criminals and how often it convicts innocent radios. We will then spend the rest of the paper discussing the appropriate punishment.

To catch an offending human, the crime must be noticed, then a suspect must be identified, apprehended, and convicted by police and prosecutors.

For cognitive radios, the crime is interference to a primary user. There are two entities who could fill the role of catching an offending radio: the primary itself or a regulator. The primary will not notice non-harmful interference, and so should participate in catching if a 'no harm, no foul' property is desired. However, it may also want to 'cry wolf,' so care must be taken to keep the rate of innocent conviction low.

For the regulator to catch offending radios, it must deploy a dense sensor net to keep track of transmission. However, it would not be able to tell the difference between innocuous and harmful transmissions. The required infrastructure is not necessarily a large added expense: [11] discusses a sensing as a service architecture that could be leveraged for enforcement purposes as well.

Identifying the offending radio is already an FCC requirement for whitespace devices [12] and is an interesting and complex question in its own right. [13], [14] look at some of the considerations for creating an identity system for cognitive devices. In particular, using the time-code idea (versus a header-implementation) in [13] may allow simpler distinguishing of interference and random noise. The actual bit-sequence of the codes will also affect the ability to catch the right offender: classically [15], identity codes with good performance force any one bit-sequence to apply to multiple different radios. This causes innocent radio to be accused if their code is the same as an offending neighbor.

Finally, with a suspect in mind, a radio must be accused and potentially convicted. For simplicity we assume strict liability, and so any radio that is accused of causing interference will be immediately punished. The accusation is a simple 'go-to-jail' command sent by either regulator or primary, depending on the implementation, that the radio is certified to follow. Notice that this is all certification must accomplish, so it is in fact easy to do and applicable to all possible devices.

We will treat the system that catches the offending (or innocent) radio as a black box with performance metrics of how well it catches a guilty radio and how often it convicts

an innocent radio. This is because before trying to design the entire architecture, we must first ask whether punishing cognitive radios will actually be effective in deterring them from causing harmful interference.

For humans, the question of deterring rational criminal behavior with appropriate punishments is covered by a large body of work on the economics of law. Seminal works such as [16], [17] see the criminal as a rational player in a game. To see applications to several areas of law, see [18], [19].

In work on the individual rational criminal, it is assumed that all humans value at least two universal quantities: money and leisure time. The extent to which each is valued depends on your socio-economic status, and therefore the punishment must change appropriately [20]–[22]. But any punishment that includes a sufficient quantity of one or both will be effective at deterring rational, human criminals.

However, a 'rational' cognitive device that responds appropriately to specific kinds of 'pain' does not make as much sense. Radios cannot feel pain or value things like money or free time. But whereas we cannot certify people, we can and should certify radios. So, we can engineer preferences.

Issuing sanctions against a radio actually aims their effect against the decision-making designer, service provider, or user involved. But who is responsible for a misbehaving radio? If we want to let the general public use devices without having to issue the equivalent of drivers licenses, the device should be built such that it is hard for a user to innocently cause it to misbehave. Between the designer or service provider, the liable party may depend on the device.

We proposed in [1], [2] a simple solution which avoids the question of which person is liable by engineering the device to degrade its own performance when it is caught cheating. This way, all three groups of people are affected. If the individual user is at fault, their device becomes less reliable. If the manufacturer or service provider have created a deviant device, it will become known to not be reliable and therefore will lose market appeal.

In [1], [2], we addressed a kind of worst-case radio that was throughput-greedy. So, we created a jail sanction that degraded performance by forcing a radio to stop transmitting for a while when caught cheating. However, not all radio applications care only about maximum data rate. A node in a sensor network may have only one packet to send every few days. Waiting another five minutes is irrelevant. On the other hand, it may be very difficult to replace the batteries on this sensor node. Even if the node was able to harvest energy from the environment [23], it will have a maximum rate of energy consumption.

These energy users seem to be another corner-case. So in this paper, we target them with a 'singing' sanction added to our original jail model. This means that a device must transmit gibberish in a garbage band while sitting in jail. However, the sanction could also be implemented with a resistor that burns energy as heat. Or, it could involve sensing for a primary and transmitting the results as community service. It could even be used as a vehicle for allowing a social sanction: if a whole network is collectively punished for interference, being forced to tell your neighbors your name when caught will give them the opportunity to punish you further.

B. What we found

Just as all people care at least about money and leisure time, the real radio ecosystem consists of radios that will at least care about some combination of data rate and energy. In Section II, we show that even having pure-time and pure-energy users coexisting with the same jail-only sanction causes significant degradation to the time-user's performance. Therefore, a singing sanction is required.

In Section III, we look at how harsh this singing sanction must be to enforce sharing rules for all pure-energy users. We see that just as we required a homeband in [1], [2], we require something similar here to deter cheating for all users in all situations. In the previous papers, the homeband was a clean band where the cognitive user had primary privileges; it was used to make cheating less attractive because you would lose your homeband whenever you were caught. Here, the homeband fills a different role: when it is too expensive to operate honestly in the cognitive band (either because it costs too much to sense, or the primary is active almost all the time), the homeband is a less-costly alternative. Users that should not be operating in the cognitive band anyway have somewhere else to go.

In Section IV, we question the assumption that pure-time and pure-energy users are the actual corner cases by looking at uses that care about *both*. We see that in fact in some cases users with a particular convex combination of time and energy costs are harder to deter from cheating than the pure users. So, in this section we find a sanction that is sufficient to deter cheating for convex combinations of preferences. We also introduce a regret-based optimization to choose the particular best combination of jail-time and strength of singing.

Finally, we look which parameters are most important for keeping the overhead of the jail-system low. For honest users, the most important factor is the cost of sensing, which is intuitively obvious because this determines how expensive it is honestly watch for the primary. With the jail system, the most important parameter for keeping the actual sanction low is the band expansion: how large the cognitive band is relative to the size of the homeband. To keep overhead low, how well you catch actual offenders matters *much* less than the rate of convicting innocent radios [24].

II. ENERGY-SENSITIVE DETERRENCE UNDER TIME-SPECIFIC SANCTIONS

We will use the same model and assumptions as in [1]. We assume that time is split into blocks in which the cognitive user can send one packet of its message using one unit of energy regardless of whether the primary is also present. These users are a legitimate class that are constrained by bandwidth instead of power [25], and represent a worst case of not being deterred from cheating simply because the primary activity makes the cognitive band less attractive [26].

As shown in Fig. 1, we assume a primary that has a probabilistic on-off pattern characterized by P_{tx} , the probability that the primary is on, iid across time steps. This allows us to simplify the actions of the cognitive user: if the primary were not iid, then the cognitive user would want to employ a more

sophisticated action set than simply cheating or not at each time step independently.

We assume that there exists a catching mechanism with performance metrics P_{catch} , the probability that a cheating cognitive device will be caught, and P_{wrong} , the probability that an honest cognitive device will be caught wrongfully when the primary is transmitting. P_{wrong} could apply even when the primary is not active: the device could be operating in many band and have a global jail. We do not consider this case here.

We assume that cognitive devices will follow the go-to-jail command without question. We will also assume that the radio, when caught, will sit in jail without transmitting for a random time that is geometric with mean $1/(1 - P_{pen})$.

The radio acts in response to the primary: it will transmit whenever the primary is not active. It will transmit when the primary *is* active with probability P_{cheat} , chosen to minimize its cost. The cognitive device is assumed to be risk-neutral.

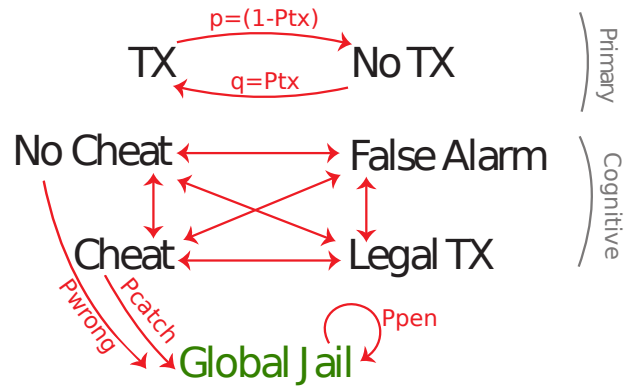


Fig. 1. The Markov chain on which the enforcement game is played. The primary transmits with probability P_{tx} , iid at every time step. The cognitive user legally transmits when the primary is active and cheats with probability P_{cheat} otherwise. The cognitive user is sent to jail with probability P_{catch} if cheating and probability P_{wrong} if not, where it must wait until a random timer determined by P_{pen} runs out.

We consider both time and energy-sensitive users. The time users are the same as those in [1], and have cost function:

$$C_T = \frac{1}{\pi_{tx}} = E[\text{time required per transmission}] \quad (1)$$

where π_{tx} is the long term average amount of time spent transmitting.

The energy-sensitive user, like a sensor-network node has a limited battery life, which it wants to spend only on actual transmissions. We define the energy cost to be the amount of energy required per transmission of actual message. If checking for the primary were free, then the cost would be $C_{E,tx} = \frac{\pi_{tx}}{\pi_{tx}} = 1$, which does not depend on the jail sentence. These users are completely ambivalent to the primary.

But checking for the primary *does* cost energy: [27] shows that the processing power required to check whether the primary is active is non-trivial. So, assuming that sensing is instantaneous, always correct, but costs some energy, we have a new utility function, $C_{E,tx,sense} = \frac{(1+C_{sense})\pi_{tx}+C_{sense}\pi_{wait}}{\pi_{tx}}$, where C_{sense} is the cost of checking for the primary and π_{wait} is the amount of time the secondary is politely waiting for the primary to turn off. We will assume that if the secondary

always cheats ($P_{cheat} = 1$), then it is never checking for the primary and so does not have to pay this sensing cost.

From an enforcement perspective, the situation is even worse: sitting in jail is free, but waiting for the primary to vacate the band costs energy. A rational user with this utility function will always cheat.

But jail does not have to be free for the energy user. If the radio must listen for a ‘get out of jail’ command, then it must remain active while in jail and therefore must burn energy. We will define this energy cost as some fraction δ of that required for sensing. δ could be very low if sensing must take place over a wide band [27], or it could be much higher if sensing is easy and the go to jail command is inside a kind of virtual puzzle box [28]. The updated cost is:

$$C_{E,tx,sense,j} = \frac{(1 + C_{sense})\pi_{tx} + C_{sense}\pi_{wait} + \delta C_{sense}\pi_j}{\pi_{tx}} \quad (2)$$

where π_j is amount of time spent sitting in jail.

By making the jail time very long, the cost of sitting in jail can rival the cost of politely waiting and therefore energy-sensitive devices can be deterred from cheating. Fig. 2 shows how high P_{pen} must be for energy-sensitive users.¹ Compare this to the P_{pen} required for time-sensitive users (green line). If the cost of jail is high relative to the cost of sensing, then the required P_{pen} can be comparable or even much lower than that required for the time-sensitive user.

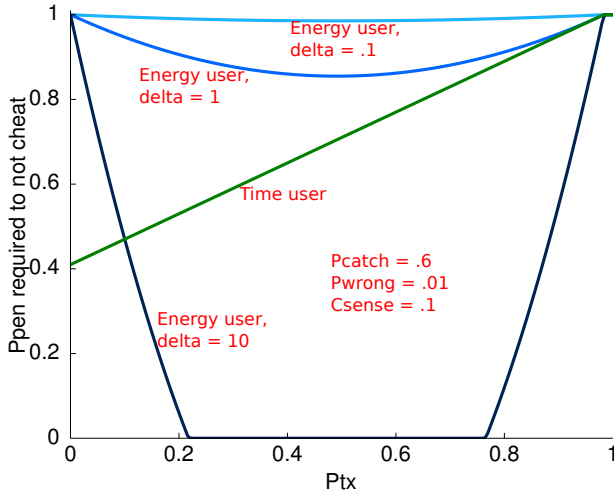


Fig. 2. The energy-sensitive user spends C_{sense} energy units checking for the primary and δC_{sense} looking for a ‘get out of jail’ command. This is the P_{pen} required to make the rational energy-sensitive user behave correctly. If the primary is almost always active or almost never active, checking for the primary does not provide much information. So, P_{pen} must be very high to make the cognitive user check anyway. Compare these values to the P_{pen} required for the time-sensitive user (green line).

Both pure users and users that care about both time and energy will likely coexist in a real wireless ecosystem. If we assume that the pure users are the worst cases, and need to account only for those two at the same time, we must use the maximum P_{pen} amongst the two types. Unfortunately, if one requires a significantly larger P_{pen} than the other, the second will suffer an overhead cost.

¹The code to produce all figures can be found here [29]

In Fig. 3, we show the overhead that the time-sensitive user incurs from having to operate under the maximum P_{pen} for the two types of users. When the cost of sitting in jail is very small relative to the cost of sensing, the time-sensitive user loses a significant fraction of the utility it would have gotten under the correct sanction. It is losing this utility from P_{wrong} : the time-sensitive user is not cheating, but it is still sent to jail erroneously with some probability.

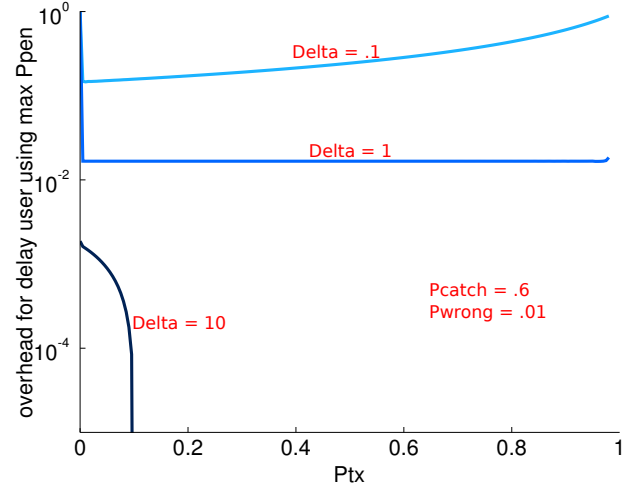


Fig. 3. To deter both time and energy users, we must use the worst case P_{pen} . This is the overhead the time-sensitive user experiences because of having to coexist with energy-sensitive users. In order to control both kinds of users without severely degrading one, we propose an energy-specific sanction in which radios must sing – transmit something not personally valuable – while sitting in jail.

Notice, however, that if the cost of jail is high relative to the cost of sensing, then the time user is affected only when P_{tx} is very small and even then not by much because the wrongful conviction rate is very low in this regime. This suggests that if we could *engineer* a high δ , then perhaps universality over energy and time-sensitive users would not cause excess burden on either.

So, we propose a singing sanction C_{sing} , which will be the energy cost of sitting in jail. For the purposes of this paper, the calculations are ambivalent to the form of the energy sanction. We just assume that the energy used can be measured relative to the energy required to transmit a message in the primary band.

With the new singing sanction, the cost for the energy-sensitive user is:

$$C_E = \frac{(1 + C_{sense})\pi_{tx} + C_{sense}\pi_{wait} + C_{sing}\pi_j}{\pi_{tx}} \quad (3)$$

III. SINGING SANCTION

A. Does it work?

As described in the Introduction, our goals are that the enforcement system will be easy to certify, universal over all radios and operating conditions, and not impose too much extra burden on an honest user within the system. Our proposed system is easy to certify because a simple lab test

can be performed to check that a radio responds to a go-to-jail command and sings for an appropriate length of time. Throughout the rest of the paper, we will rate the performance of the jail system against the ideal performance of an honest device that always senses and is not subject to jail. In this section we will look at whether it is possible to set a singing sanction that is sufficient to deter cheating by any energy-sensitive device with any cost of sensing, and any primary activity characteristic.

We need one further assumption: it is reasonable for energy-sensitive devices to have one message to send and then want to turn off for a very long time. During this long off-period, if the cognitive device is subject to wrongful convictions, its cost in (3) will go to 1 regardless of its cheating behavior. So, in order to make deterrence even possible, these radios must be able to declare themselves off and not be subject to jail.

We note that the cognitive user will always choose $P_{cheat} = 0$ or 1. This will not be proven here; the proof can be found in [30]. Therefore, we can calculate the required C_{sing} to deter cheating of cognitive users by comparing the expected long-term cost when $P_{cheat} = 0$ to that when $P_{cheat} = 1$.

We find the conditions on the enforcement parameters so that no rational user will cheat to be [30]:

$$\frac{C_{sing}}{1 - P_{pen}} > \frac{C_{sense}}{P_{tx}((1 - P_{tx})P_{catch} - P_{wrong})} \quad (4)$$

The term $C_{sing}/(1 - P_{pen})$ makes sense as a metric because it is the average cost of one stay in jail. So, the singing sanction will work if set for specific values of C_{sense} and P_{tx} , but can we set C_{sing} high enough such that *no* value of C_{sense} or P_{tx} will cause the user to want to cheat?

Fig. 4 shows the required singing sanction $C_{sing}/(1 - P_{pen})$ as a function of the cost of sensing. Notice that it scales linearly. So, we cannot deter all types of users with the same sanction.

Notice also in Fig. 4 how the required sanction changes with different P_{tx} . This suggests that some kinds of users are more difficult to deter, but so are some types of situations. In Fig. 5, we show this tradeoff directly. It is difficult to enforce no cheating when the primary is almost never present (low P_{tx}) or almost always present (high P_{tx}).

The two problem areas come from different sources. When the primary is rarely active, enforcement is difficult because when a radio is honest, it must sense all the time and will be sent to jail every once in a long while. However, when it is dishonest, it never has to sense and still goes to jail only rarely. It will cheat if it is more expensive to sense than to go to jail whenever the primary actually shows up.

High P_{tx} is a problem because of P_{wrong} . Consider an example in which $P_{catch} = 1$ and $C_{sense} = 0$. When the cognitive user cheats, every transmission requires the cost of transmission and one stay in jail. If the cognitive user is honest, every transmission requires the cost of transmission and the expected number of stays in jail. It is not possible to deter cheating if an honest user would be mistakenly sent to jail on average more than once per transmission.

We encountered the same difficulty with P_{tx} for time-sensitive users in the previous papers, and solved the problem

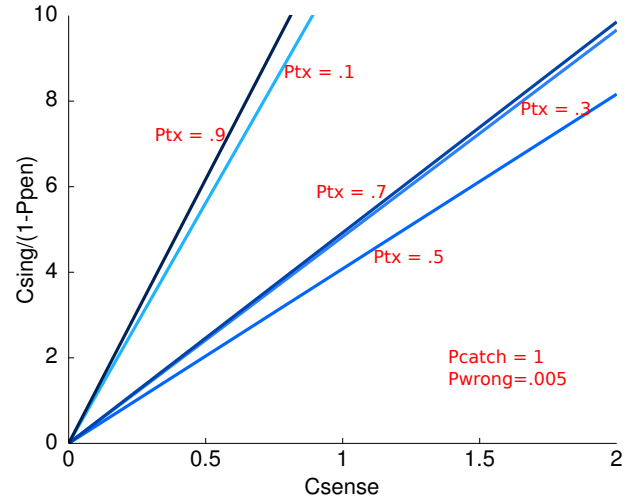


Fig. 4. The amount of singing required to make an energy sensitive user not cheat. Notice that it scales linearly with the amount of energy required to sense. So, it is not possible to set one value of C_{sing} that will be sufficient to deter cheating for all values of C_{sense} .

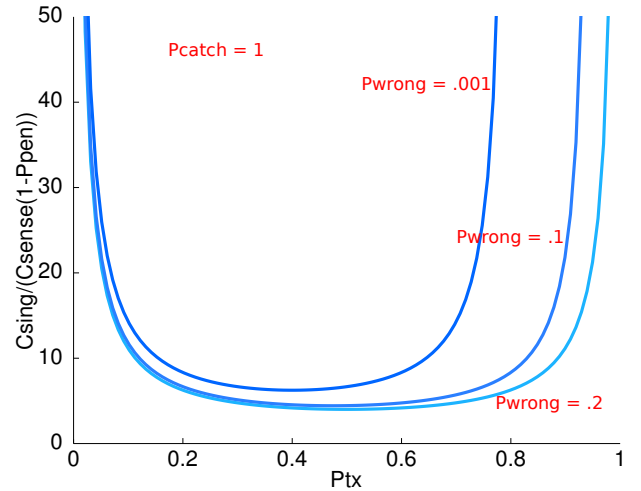


Fig. 5. The amount of singing required to make an energy sensitive user not cheat by P_{tx} . If the primary is almost never active, sensing provides little information, and jail is an infrequent event. So C_{sing} must be very high. If the primary is almost always active, it costs less to go to jail once per transmission than to wait for a free slot and be sent to jail wrongfully several times on average per transmission.

with a homeband. This is a band of size B , relative to the cognitive band, in which the cognitive user has primary privileges. The cognitive user must stake its homeband against unlawful cognitive use by being unable to transmit in either band whenever it is sent to jail. In the next section, we consider whether a similar idea will work for the energy-sensitive users.

B. Introducing a dirty homeband

If the homeband has the same cost as the cognitive band, an energy-sensitive user will have cost 1 if they use the homeband and cost greater than 1 in the cognitive band. Device manufacturers choose where they want their devices to operate. If they know that the device will get strictly better performance in the homeband alone, the device will be designed and certified to only run in the homeband. We

want the devices to fill spectrum holes, so the clean homeband solution is overkill.

We propose instead introducing a dirty homeband of energy cost C_{hb} per unit of transmission. C_{hb} is controlled by regulators by offering protection to the primary at a particular level, allowing a few cognitive devices to share a homeband, or even using an unlicensed band with a well-designed etiquette. We also propose that devices must be able to choose their operating band at *runtime*. This way, intuitively, if the cognitive band is more attractive, the device will fill spectrum holes there. However, if the band is not attractive (if the primary is always active, or if the cost of sensing is just too high), the device can choose to operate only in the homeband. This would allow the holes to be filled only when it is reasonable to do so.

With the homeband, we get the usage profile Fig. 6 for different costs of sensing and singing. The brown dashed line separates the cost of sensing when an honest cognitive user will choose to operate in just the cognitive band or just the homeband.

Two interesting effects show up with this dirty homeband. First, when C_{sing} is high enough for the cognitive user to not cheat in the cognitive band, the cognitive user subject to jail will switch to the homeband at slightly lower C_{sense} than the honest user. In this case, the homeband is acting as a safe haven; the cognitive user switches to the homeband to be free from wrongful convictions.

The second effect is for high C_{sense} . At some value of C_{sing} , the cognitive user will choose to operate in the homeband instead of cheating in the cognitive band. This is similar to an effect in the human realm with hackers who are illegally hacking for profit [31]. If they have the opportunity to use these same skills for a legitimate purpose (as in a security company), then sanctions need to be just high enough to make the legitimate job tempting. They do not need to be high enough to deter these hackers from illegal activity. Another example is the famous con-artist Frank Abagnale Jr. [32]. As a con-man, he was completely undeterrable. However, when offered a legitimate security-consultant position with the FBI, he gave up a life of crime for a far more lucrative career. The arms-race to keep Abagnal in line was not necessary when a suitable legal alternative was available.

To formalize the cost structure of energy users with a homeband, we give the costs of each of its rational options:

- 1) Transmission in only the cognitive band:

$$C_{E,1} = \frac{(1 - C_{sense})\pi_{tx} + C_{sense}\pi_{wait} + C_{sing}\pi_j}{\pi_{tx}} \quad (5)$$

- 2) Transmission in only the homeband:

$$C_{E,2} = C_{hb} \quad (6)$$

- 3) Transmission in *both* the home and cognitive bands:

$$C_{E,3} = \frac{(1 + C_{hb}B + C_{sense})\pi_{tx} + (C_{hb}B + C_{sense})\pi_{wait} + C_{sing}\pi_j}{B\pi_{nj} + \pi_{tx}} \quad (7)$$

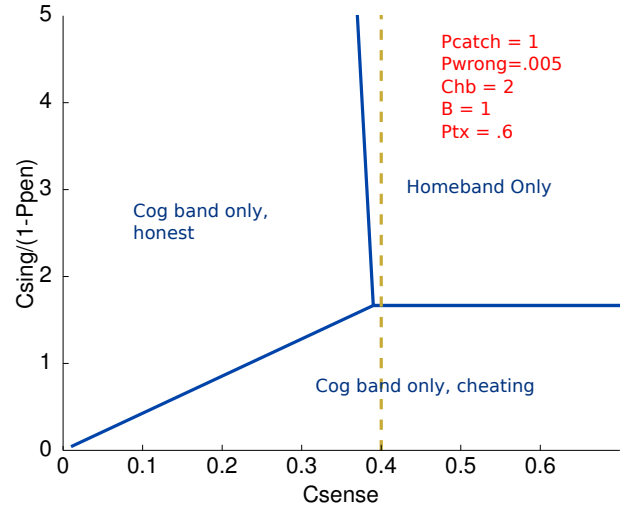


Fig. 6. The usage profile of a user with sensing cost C_{sense} for different levels of C_{sing} . A pure energy user will choose between cheating in the cognitive band, honestly playing in the cognitive band, and operating only in the homeband. The dashed line separates where an honest user not subject to jail will play in the cognitive band (for low C_{sense}) or in just the homeband. We see the dual roles of the homeband here: when the jail-enforced user switches at a lower level of C_{sense} than the honest user, the homeband is acting as a safe haven from wrongful convictions. With high C_{sense} and high C_{sing} , the cognitive user operates only in the homeband because it is a legitimate, attractive alternative to cheating in the cognitive band. Notice that there exists a value of C_{sing} that will be sufficient to deter cheating for all users simultaneously.

So the overall cost of the energy user is:

$$C_E = \min(C_{E,1}, C_{E,2}, C_{E,3}) \quad (8)$$

Note that the rational energy-sensitive user will never choose to operate in both bands because it does not care about how quickly its messages are sent.

Fig. 7 shows the usage profile for P_{tx} vs C_{sing} . The dotted lines separate the values of C_{sing} where a rational user will choose to cheat or not in the cognitive band. The solid lines separate the regions where a rational user will operate in the homeband alone vs the cognitive band. In order to deter cheating at any particular P_{tx} , C_{sing} must be set to the minimum of these two lines.

Notice first that high P_{tx} is no longer a problem with the introduction of the homeband. If the primary is active too much of the time, the cognitive user will simply switch to the homeband.

If the primary is rarely active, however, it is still very difficult to deter the radios from cheating. So, we must ask whether we actually *want* to deter these users from cheating. This is a policy question that really depends on the kind of primary. If the primary rarely uses its allotted band, perhaps it is better to force this primary to operate as a cognitive device elsewhere. Perhaps we want to accept a kind of ‘use it or lose it’ rule for these primaries or even allow cognitive users squatters’ rights. Our sanction then must be set to cover any primaries with a high enough level of usage.

The primary could alternatively look like a public-safety service. These users rarely are active, but when they are active, they have high usage for a period of time and *must*

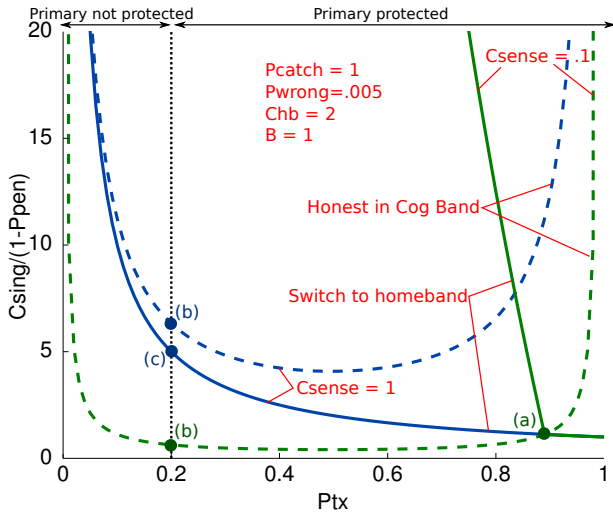


Fig. 7. How a cognitive device will behave by P_{tx} . The dashed lines separate regions where the cognitive user will not cheat (above) and cheat in the cognitive band. The solid lines separate regions where the cognitive user will choose to operate only in the homeband (above the line). Blue lines have $C_{sense} = 1$ and green lines have $C_{sense} = .1$. The black dotted line shows our choice of the lowest P_{tx} for which we will design C_{sing} to protect the primary. In order to choose a sufficient C_{sing} for all P_{tx} , there are three points that must be considered: (a) where the dashed and solid lines cross, (b) where the no cheating line crosses the lowest guaranteed P_{tx} , and (c) where the homeband-switching line crosses the lowest guaranteed P_{tx} . Letting $C_{sing}/(1 - P_{pen}) = \max(a, \min(b, c))$ gives the required sanction.

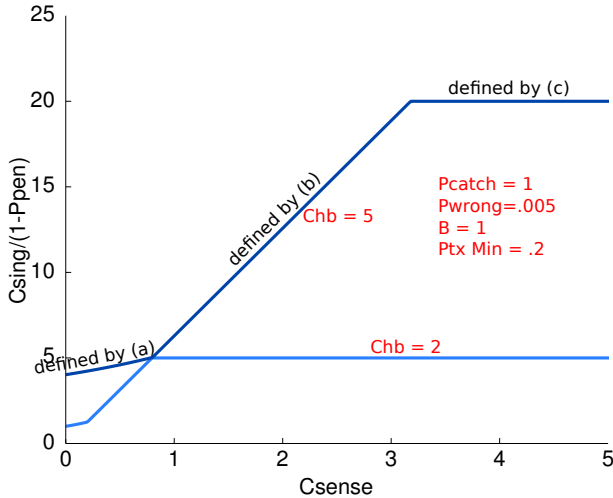


Fig. 8. If the sanction is chosen to be universal over all $P_{tx} > P_{tx, min}$, this is the required sanction by cost of sensing. The three regions of the plot are where different points in Fig. 7 are determining the required sanction.

be protected. Having legitimate services that have periods of low and high use calls into question our assumption of an iid primary. Perhaps for these users a different kind of protection guarantee must be implemented. From [33], we know that realistically a cognitive user must have some buffer zone in which it is allowed to interfere before it realizes the primary is active. We may not want to force cognitive users to sense if the primary is almost never active. But, when the primary starts a high-usage period, cognitive users start being sent to jail more often, and will realize the primary is once again active. So, our sanction again can be set to protect the primary during

a high-usage period assuming the cognitive devices will do whatever they want when the primary is very rarely around.

The final situation is one in which we *must* protect these low usage primaries. In order to keep overhead reasonable, the options are to disallow cognitive use or implement an adaptive jail that changes based on how often a user is caught cheating.

For this paper, we will assume that we have iid primaries that do not change their activity pattern. Also, we will assume that there is a lower bound on the P_{tx} we want to protect. All other options are left for future work. In Fig. 7, this minimum P_{tx} is set at 0.2. If we want a sanction that is universal over the P_{tx} above the minimum, we need to set $C_{sing} = \max(a, \min(b, c))$ where (a), (b), and (c) are the indicated points in Fig. 7.

Fig. 8 shows the required C_{sing} for each C_{sense} such that the sanction works for all P_{tx} greater than the minimum. Again, we see that the required C_{sing} saturates, so we can choose a C_{sing} that will work universally over all sensing technologies and all primary transmission characteristics.

Using this universal C_{sing} , we have a new set of users and conditions that will choose to play in the cognitive band versus switch to the homeband. This usage profile over P_{tx} and C_{sense} is shown in Fig. 9. Below the minimum P_{tx} , some users will still choose not to cheat, but there is no guarantee that those with higher sensing costs will be honest. Again, the honest user without jail is shown with the brown line.

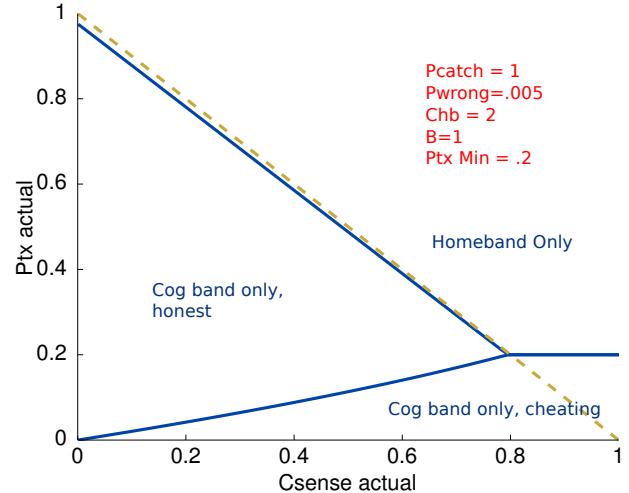


Fig. 9. Using the maximum required sanction over C_{sense} from Fig. 8, this is the actual usage profile for users with sensing cost C_{sense} and primary characteristic P_{tx} . We do not guarantee the cognitive user will not cheat below $P_{tx, min} = .2$.

Fig. 10 shows the overhead associated with sensing and the jail system relative to a genie user who knows for free when the primary is transmitting. Most of the overhead is coming from having to sense, as shown by the overhead for the honest user with no jail. Jail sanctions that are more specific to the individual user naturally have better overheads. Finally, the overhead is capped by the relative cost of the homeband to the primary band because the genie user will always stay in the cognitive band while any that have to sense will switch to the homeband for high enough sensing costs.

The overhead coming from jail is caused by wrongful

convictions. This can be seen in Fig. 11 when we look at the overhead after changing P_{wrong} to 0.05. Now, users vacate the cognitive band for lower levels of C_{sense} . They also have significantly more overhead because they must spend much more time wrongfully in jail.

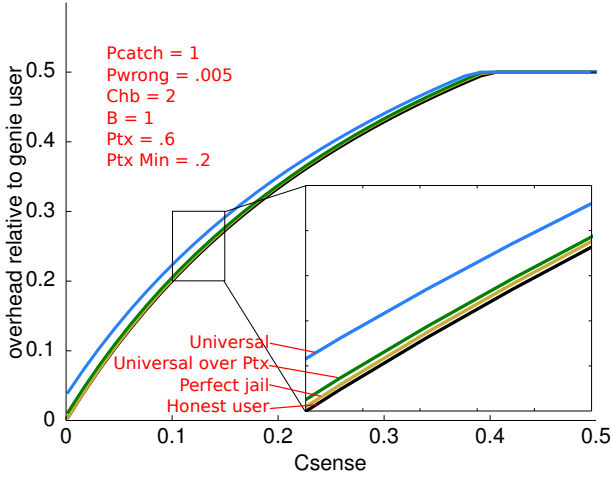


Fig. 10. The overhead of users in different kinds of jail systems relative to a genie user who always knows for free when the primary is active.

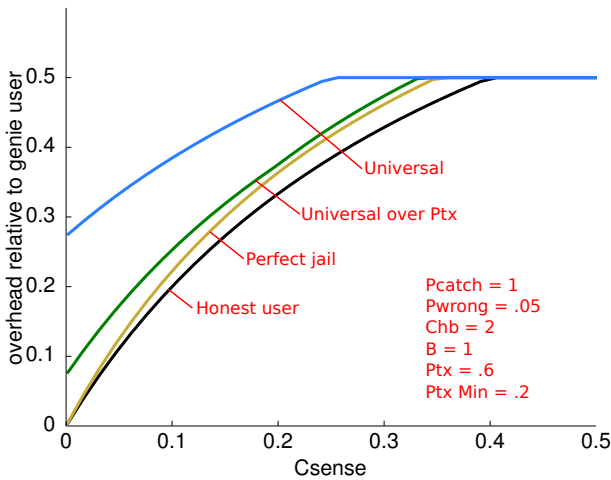


Fig. 11. The overhead of different universalities of jail relative to a genie user for a higher P_{wrong} . P_{wrong} is a strong determining factor in the overhead incurred because it is responsible for all of the energy spent in jail. The effect of P_{wrong} is particularly bad when the cost of sensing is low.

IV. MIXED ENERGY AND TIME-SENSITIVE USERS

A. How it works

We can deter cheating in both time and energy users by setting the length of the jail sentence for the time-sensitive users, and then setting the added singing penalty for the energy-sensitive users.

However, realistic devices will not always care about only time or only energy. Consider a smartphone downloading pictures. Download speed matters for the user. But if only one picture can be downloaded before the phone runs out of batteries, it is not a very useful mobile device. For these users

that care about both time and energy, will setting sanctions based on the two extremes be sufficient?

There are many ways to define a total cost for devices that care about both energy and time. For simplicity, we consider only convex combinations of pure energy and pure time costs:

$$C_{mixed} = (1 - k)C_T + kC_E \quad (9)$$

We will refer to these as mixed devices. We leave for future work understanding how the results here generalize to different total cost formulations.

In order to understand how to set the appropriate jail length and strength of the singing sanction, we first must notice that these parameters no longer trade off as $C_{sing}/(1 - P_{pen})$. This is because the time-cost is sensitive to only the length of the jail sentence, while the energy-cost is sensitive to both parameters. So, to understand how these enforcement parameters work together, we will first assume that P_{pen} is a fixed value and will investigate how C_{sing} must be set to deter cheating. Then later we will look into how to optimally choose both values.

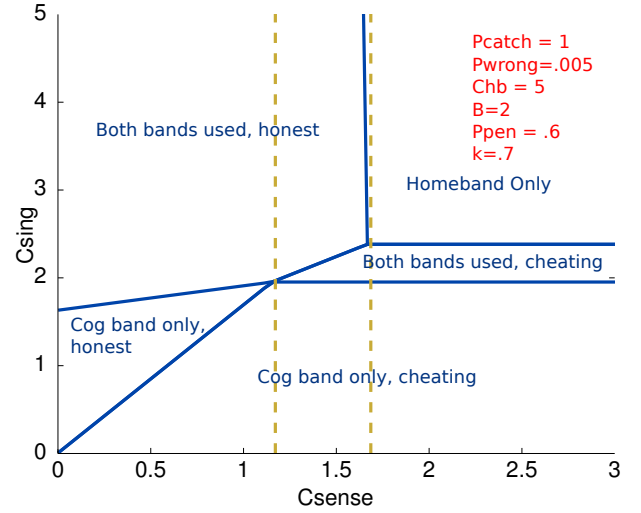


Fig. 12. The usage profile for a cognitive device that cares about both time and energy. Now, the desire for higher rate outweighs the desire to transmit for the best energy performance in some cases. Therefore, this mixed user will have regions in which it uses just the cognitive band, both cognitive and home bands, or just the homeband alone. The brown, dashed lines indicate where an honest user with no jail will use the cognitive band alone (low C_{sense}), both bands (mid C_{sense}) or just the homeband (high C_{sense}).

Fig. 12 shows the response of a user with $k = 0.7$ and different costs of sensing to different levels of the singing sanction, C_{sing} . Unlike for the purely energy-sensitive users, the mixed users will at times operate in just the cognitive band, both the cognitive and home bands, or only in the home band. But like the pure-energy case, there seems to exist a level of C_{sing} that is sufficient to deter cheating for *any* cost of sensing C_{sense} .

The energy costs were given in the last section. The time costs are formalized as:

- 1) Transmission in only the cognitive band:

$$C_{T,1} = \frac{1}{\pi_{tx}} \quad (10)$$

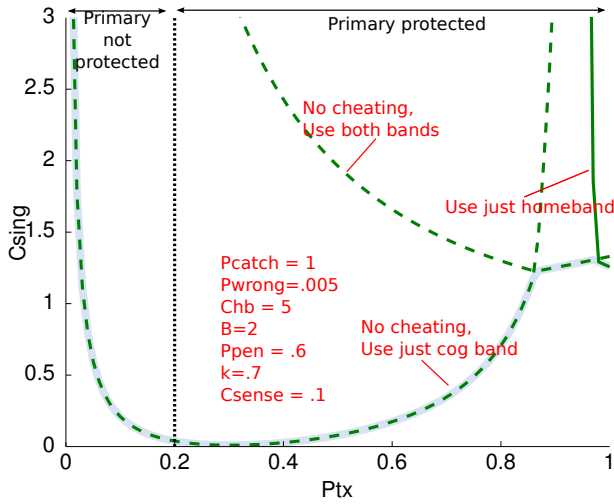


Fig. 13. The usage profile by P_{tx} . For any mixed user, there are regions in which it will cheat or not cheat using just the cognitive band (lower dashed line), both bands (higher dashed line), or just the homeband (solid line). There are again the same three points used to determine the required C_{sing} over all $P_{tx} > P_{tx,min}$. However, now the maximum of these must be taken over all k to get a sanction that deters cheating with all types of users.

2) Transmission in only the homeband:

$$C_{T,2} = \frac{1}{B} \quad (11)$$

3) Transmission in *both* the home and cognitive bands:

$$C_{T,3} = \frac{1}{B\pi_{nj} + \pi_{tx}} \quad (12)$$

So the overall cost of for the mixed user is:

$$C_{mixed} = \min_i((1-k)C_{T,i} + kC_{E,i}) \quad (13)$$

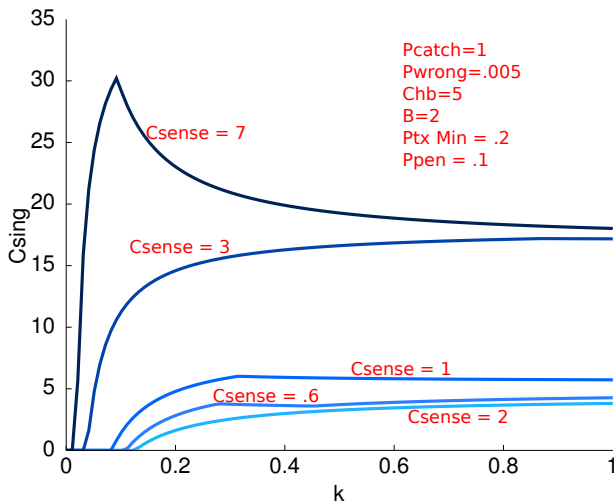


Fig. 14. The sanction required to cover all protected P_{tx} for different values of k and C_{sense} . When P_{pen} is high (not shown), the worst case k is the pure-energy user. If P_{pen} is low, as shown here, when C_{sense} gets large, the worst case k is some mix between time and energy users. So, in order to get a sanction that is sufficient to deter all mixed users, we must look at the maximum required over all k .

We want to set C_{sing} to deter cheating universally over all primary transmission characteristics as well, so Fig. 13

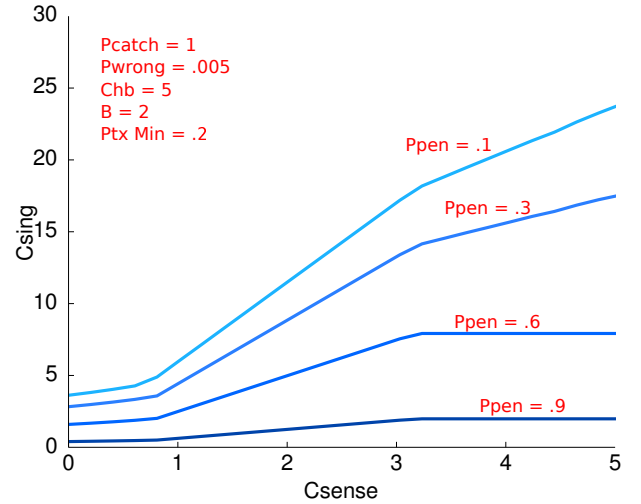


Fig. 15. Taking the sanction required to be sufficient for all P_{tx} and k , this is what is required over C_{sense} for different lengths of time of jail. For high P_{pen} , the required C_{sing} levels out when the worst case of the pure-energy user is satisfied. For low P_{pen} , the required C_{sing} will eventually level out, but at a much higher C_{sing} with a much higher worst case C_{sense} .

shows the behavior of the cognitive user seeing different primary transmission levels P_{tx} at different levels of the singing sanction. Notice that there are three regions that we care about: when C_{sing} is high enough to deter cheating by users operating only in the cognitive band, deter cheating by users operating in both bands, and causing users to use only the homeband. We once again have the problem of low primary presence making enforcement very difficult. So, we will again assume there is a minimum $P_{tx} = 0.2$ below which we do not guarantee the sanction will deter rational users from cheating. For all other values of P_{tx} , we can set C_{sing} to the minimum of what is required for the three different use cases.

Fig. 14 shows the C_{sing} required to be universal over P_{tx} against the value of k , the parameter determining the preference between time and energy cost. We show in this plot only the surprising case: when P_{pen} is very small. If P_{pen} is large, the required C_{sing} is dominated by the energy user for all costs of sensing. However, if the length of jail is small, Fig. 14 shows that the worst-case k is not one of the extremes. We must take the maximum sanction over all k .

Fig. 15 shows the sanction required to deter bad behavior for any k and any P_{tx} given a cost of sensing. As before, there exists a C_{sing} such that all users will be deterred from cheating regardless of preference, sensing cost, or primary usage characteristic. The sufficient values of C_{sing} and P_{pen} are shown in Fig. 16. Any value of the enforcement parameters above the line will be sufficient to deter cheating over all users. Notice that for high P_{pen} , the required singing value is set by the pure-energy users. For lower P_{pen} , a mixture is the worst case. If P_{pen} gets so low that it cannot deter the pure time user, the required C_{sing} will go to infinity (for the example shown here, this minimum P_{pen} is so small that it is not on the plot).

We have a tradeoff, but for a real system, how should we choose the correct value of C_{sing} and P_{pen} ? What does the 'correct' value even mean?

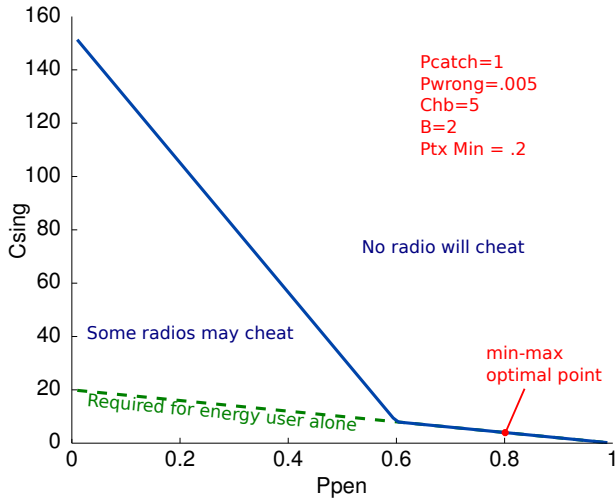


Fig. 16. The required tradeoff between C_{sing} and P_{pen} to guarantee that no radio will cheat regardless of P_{tx} , k , or C_{sense} . The high- P_{pen} range is determined by the pure-energy user. Other points are determined by a mixed user. We now have freedom to choose where on this line we want to operate. We propose that the correct point is that which minimizes the worst-case overhead over all types of users. This min-max optimal point is indicated.

B. Choosing the correct sanction

We propose that the parameters should be set by a regret-based optimization, similar to the ideas in [34]: the fact that we are choosing a sanction that will work for any type of user facing any type of primary implies that some users will incur a higher cost (overhead) than if the sanction were tailored specifically for them. We propose to choose the actual sanction to minimize the worst case overhead over types of users. Formally, we define overhead as the utility lost relative to an honest user with no jail:

$$overhead = 1 - \frac{1/C_{mixed}}{1/C_H} \quad (14)$$

where C_H is the cost of an honest user. Then, we choose the sanction according to:

$$\min_{(C_{sing}, P_{pen}) \in \Phi} \max_{k, C_{sense}} overhead \quad (15)$$

where Φ is the set of all pairs (C_{sing}, P_{pen}) that will deter cheating universally over types of users and primary conditions.

The optimal min-max sanction is indicated on Fig. 16, and happens in the region defined by the pure-energy user. Fig. 17 shows the behavior of a user with $k = .7$ over different C_{sense} and P_{tx} subject to this optimal sanction. Also shown is the behavior of an honest user not subject to jail. For P_{tx} above the minimum protected value, the discrepancies between the jail-affected user and the honest user are again due to wrongful convictions making the cognitive band more expensive relative to the homeband.

Fig. 18 shows the overhead of an example user under different sanctions. The honest user is not subject to jail. The universal sanction is min-max-optimal for all users. All other jail sanctions perform somewhere in between these two cases.

We see that the majority of the overhead comes from the cost of sensing. The extra overhead for different jail

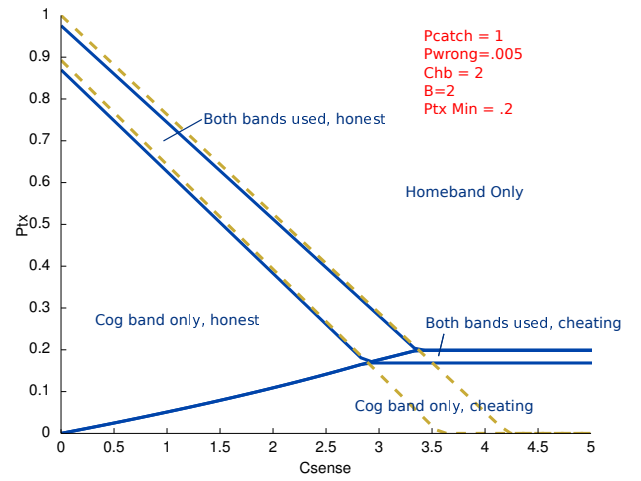


Fig. 17. Using the min-max optimal sanction, we have a usage profile for a given k and different C_{sense} .

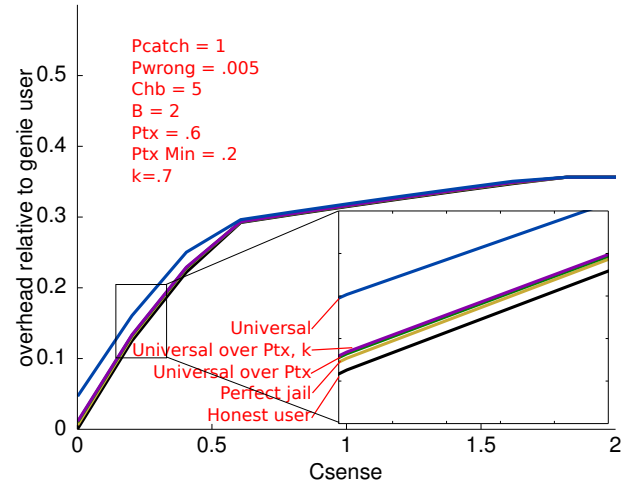


Fig. 18. The overhead of a users with different types of sanctions relative to a genie user that knows for free when the primary is transmitting.

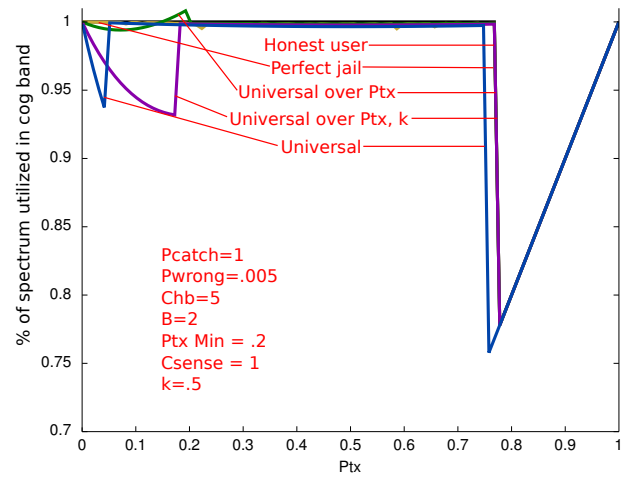


Fig. 19. The society cares about how much of the cognitive band is being utilized. This is shown here for different kinds of jail. When the primary is usually around, the cognitive devices vacate the band in favor of the homeband. The difference in P_{tx} where those with a universal jail vacate before the honest users is a kind of societal overhead of the jail system. When P_{tx} is very low, some users start to cheat; their time in jail causes holes to be left unfilled.

sanctions is coming only from wrongful convictions, with the universal sanction producing the most. So, to reduce overhead, improving sensing technology is the most important factor.

From a regulator's perspective, minimizing the overhead for a particular user is not enough. The regulator also cares about overall spectrum utilization, or how well the cognitive devices are filling the holes left in primary spectrum usage. In Fig. 18, we already see that when C_{sense} is high, even honest, jail-free devices will vacate the cognitive band.

We look directly at spectrum utilization in Fig. 19. There are three regions to notice: for middle ranges of P_{tx} , all kinds of jail are very effective at allowing the cognitive users to fill spectrum holes. For high P_{tx} , cognitive devices will start leaving the cognitive band in favor of transmitting only in the homeband. An honest user with no jail will rationally switch at some point because of the sensing cost. Users with jail sanctions will switch earlier. The regions in which the cognitive devices with jail have switched but a rational honest user is still occupying the cognitive band are a type of overhead: they are spectrum holes that remain unfilled because of wrongful convictions in the jail system.

Finally, consider low P_{tx} . Our universal sanctions do not guarantee that cheating is deterred here over all users. So, some users will begin to cheat. Because the cheating cognitive devices spend so much time in jail, the overall spectrum occupancy goes down.

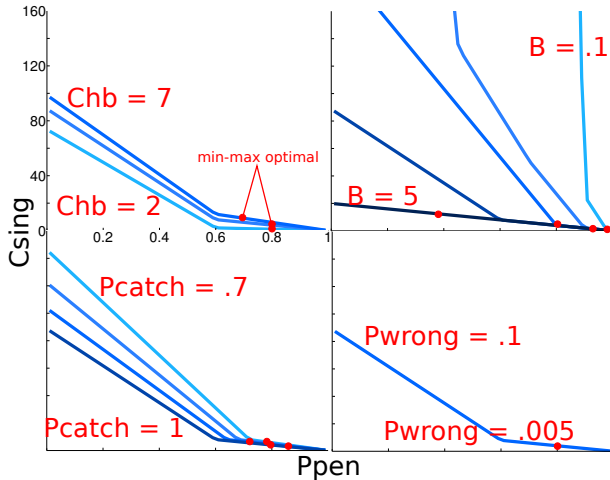


Fig. 20. To understand how our outside parameters of C_{hb} , B , P_{wrong} , and P_{catch} affect the operation of the jail system, we look at first how the universal sanction changes with these parameters. Since the min-max optimal point occurs when P_{pen} is high, the only parameter that affects the sanction much is the size of the homeband B .

Thus far, we have assumed that certain parameters are held fixed because they are determined by elements outside the jail system. These parameters are P_{catch} , the probability of being caught when cheating, P_{wrong} , the probability of being caught when not cheating, C_{hb} , the energy cost of transmitting one packet in the homeband, and B , the size of the homeband. In Figs. 20, 21, and 22, we look at the effect of these parameters on the sanction, the overhead, and the spectrum utilization to see which are important.

For the sanction in Fig. 20, notice that the only parameter with a large effect is the size of the homeband. The min-max

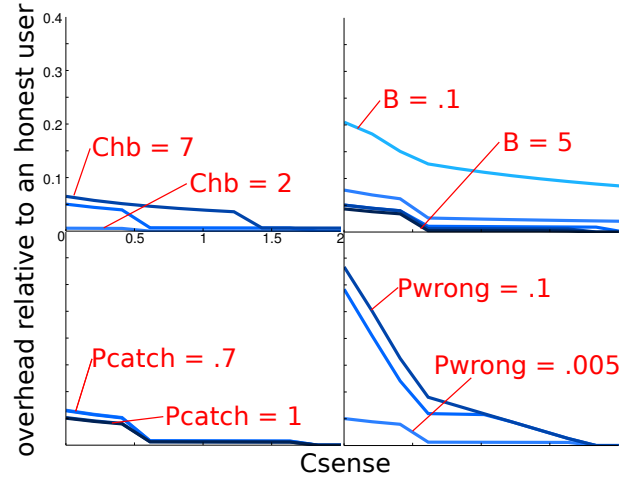


Fig. 21. The effect of changing parameters on the overhead of an individual radio relative to an honest user. B affects the overhead because of the much higher sanction. P_{wrong} is determining the amount of time spent in jail, so it is the biggest determiner of overhead.

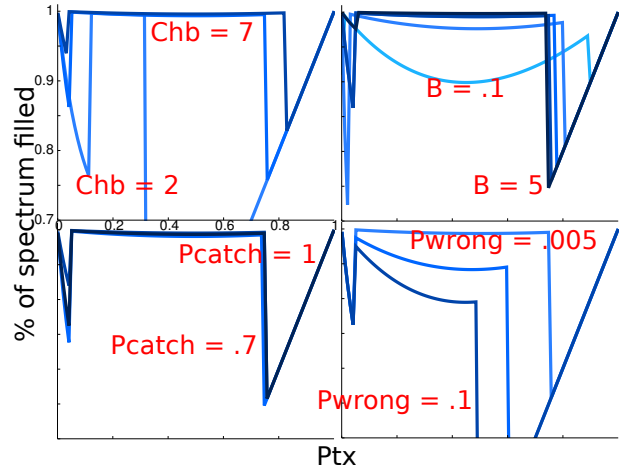


Fig. 22. The effect of changing parameters on the percentage of spectrum utilized. While users are occupying the cognitive band, the biggest determiners of the utilization are B and P_{wrong} because they are determining the how much time is spent in jail. Both these and C_{hb} also have a large effect on when the cognitive users will choose to vacate the cognitive band in favor of the homeband.

optimal point is always in the high- P_{pen} region, so changing the other parameters does not have a large effect on the actual sanction imposed. P_{wrong} in particular is interesting: the required sanction is always determined by high C_{sense} , and is not sensitive to P_{wrong} . This means that as catching technology improves, the same sanction will still be optimal.

For the overhead in Fig. 21, the size of the homeband and the rate of wrongful convictions are the two most sensitive parameters. Lower B causes a much higher sanction and therefore higher cost of wrongful conviction. The rate of wrongful convictions makes a big difference because these are the base source of the overhead of the jail system.

In Fig. 22, we see the spectrum occupancy. C_{hb} is a large factor because with a better homeband, cognitive users will choose to switch out of the cognitive band earlier than they would otherwise. B lowers occupancy again because of the

longer jail sentences. Higher P_{wrong} makes the cognitive band much less attractive because they will spend more time in jail, so users will switch to the homeband sooner.

Overall, to improve the performance of users, the cost of sensing must be made low. For other parameters, the size of the homeband is most important for keeping the actual sanction small. But P_{wrong} is the critical parameter for improving overhead. Catching the criminals is far less important than not punishing innocent users.

V. CONCLUDING REMARKS

In this paper, we have explored enforcement through a jail-based mechanism for an ecosystem of wireless devices that care about energy as well as time. We have found that this sanction satisfies our requirements for light-handed regulation: it is easy to certify, can be made universal over all types of users and technology, and it is possible to keep the extra burden to an honest user small by keeping P_{wrong} small. Having a homeband allows users that should not be using the cognitive band to transmit elsewhere, and so it makes intuitive sense that the best way to have high utilization of the cognitive bands is to make it easier for cognitive users to operate: by lowering the cost of sensing and the rate of wrongful convictions.

From here, future work includes addressing our assumptions. We have assumed an iid primary, but real primaries have periods of high use and low use. Also, we have assumed that the cognitive devices always have a message to send. This is obviously not realistic for the energy and mixed users. Finally, we have addressed only mixed users with convex combination cost functions. This is not an exhaustive characterization: cell phones exhibit a very different behavior, for example. They have a threshold requirement on the rate, and as soon as the rate is satisfied they want to minimize their energy usage. We have already seen that mixed users are in some cases harder to deter from cheating than the pure users. Will different kinds of mixed cost functions also be harder to deter from cheating?

ACKNOWLEDGEMENTS

We thank the U.S. NSF (CNS-403427, CNS-0932410 and a Graduate Research Fellowship) for their generous funding support. We also thank Pulkit Grover and the anonymous reviewers for their helpful comments.

REFERENCES

- [1] K. Woyach, A. Sahai, G. Atia, and V. Saligrama, "Crime and punishment for cognitive radios," *Allerton*, 2008.
- [2] K. Woyach, P. Pyapali, and A. Sahai, "Can we incentivize sensing in a light-handed way?," in *New Frontiers in Dynamic Spectrum, IEEE Symposium on*, (Singapore), pp. 1–12, Apr. 2010.
- [3] Federal Communications Commission, "Second memorandum opinion and order," Sept. 2010.
- [4] 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.
- [5] J. M. Peha and I. Khamitov, "Pay Cash: A Secure Efficient Internet Payment System," *Proceedings of the 5th International Conference on Electronic Commerce*, pp. 125–130, 2003.
- [6] S. Mishra, A. Sahai, and R. Broderson, "Cooperative sensing among cognitive radios," *IEEE ICC*, pp. 1658–1663, June 2006.
- [7] J. Sachs, I. Maric, and A. Goldsmith, "Cognitive cellular systems within the tv spectrum," *Proceedings of the Fourth IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Apr. 2010.
- [8] T. Hazlett, "Assigning property rights to radio spectrum users: Why did fcc license auctions take 67 years?," *Journal of Law and Economics*, p. 529, 1998.
- [9] T. W. Hazlett, "Optimal Abolition of FCC Spectrum Allocation," *Journal of Economic Perspectives*, vol. 22, no. 1, pp. 103–128, 2008.
- [10] E. Goodman, "Spectrum Rights in the Telecom to Come," *San Diego Law Review*, vol. 41, no. 269, 2004.
- [11] M. Weiss, S. Delaere, and W. Lehr, "Sensing as a service: An exploration into practical implementations of dsas," in *New Frontiers in Dynamic Spectrum, IEEE Symposium on*, (Singapore), Apr. 2010.
- [12] "FCC Adopts Rules for Unlicensed Use of Television White Spaces." http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf.
- [13] 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.
- [14] A. Sahai, K. Woyach, G. Atia, and V. Saligrama, "A technical framework for light-handed regulation of cognitive radios," *IEEE Communications Magazine*, pp. 96–102, Mar. 2009.
- [15] R. Ahlswede and G. Dueck, "Identification via channels," *IEEE Transactions on Information Theory*, pp. 15–29, 1989.
- [16] J. Bentham, *An introduction to the principles of morals and legislation, in: The Utilitarians*. Anchor Books, Garden City, NY, 1973.
- [17] G. S. Becker, "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, vol. 76, 1968.
- [18] R. A. Posner, *Economic Analysis of Law*. Little, Brown, and Company, 3 ed., 1986.
- [19] A. M. Polinsky and S. Shavell, "The theory of public enforcement of law," *Journal of Economic Literature*, vol. 38, pp. 45–76, Mar. 2000.
- [20] A. M. Polinsky, "Optimal fines and auditing when wealth is costly to observe," *International Review of Law and Economics*, vol. 26, no. 323, 2006.
- [21] 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.
- [22] L. Kaplow, "The optimal probability and magnitude of fines for acts that definitely are undesirable," *National Bureau of Economic Research Working Paper Series*, June 1989.
- [23] B. Warneke, M. Last, B. Liebowitz, and K. Pister, "Smart dust: communicating with a cubic-millimeter computer," *Computer*, p. 44, 2001.
- [24] A. Volokh, "N guilty men," *Univ. of PA Law Rev.*, no. 1, pp. 173–216, 1997.
- [25] D. Tse and P. Viswanath, *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [26] 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.
- [27] D. Cabric, M. Chen, D. Sobel, S. Wang, J. Yang, and R. Brodersen, "Novel radio architectures for UWB, 60 GHz, and cognitive wireless systems," *EURASIP Journal on Wireless Communications and Networking*, vol. 2006, no. 2, p. 22, 2006.
- [28] S. Nakamoto, "Bitcoin: a peer-to-peer electronic cash system." <http://www.bitcoin.org/bitcoin.pdf>.
- [29] http://www.eecs.berkeley.edu/~kwoyach/code/dyspan11_code.zip.
- [30] K. Woyach and A. Sahai, "Why the caged cognitive radio sings: Technical appendix." http://www.eecs.berkeley.edu/~kwoyach/papers/dyspan11_techAppendix.pdf.
- [31] P. T. Leeson and C. J. Coyne, "The economics of computer hacking," *Journal of Law, Economics, and Policy*, 2005.
- [32] S. Redding and F. W. Abagnale, *Catch me if you can: the true story of a real fake*. Broadway, 2000.
- [33] 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.
- [34] A. Sahai, K. Woyach, K. Harrison, H. Palaiyanur, and R. Tandra, "Towards a 'theory of spectrum zoning'," *Allerton*, Oct. 2009.