Outline:

George Atia, Anant Sahai and Venkatesh Saligrama,
"Spectrum Enforcement and Liability Assignment in Cognitive Radio Systems,"
IEEE DySpAN, Oct. 2008.

Paper available: www.eecs.berkeley.edu/~sahai/Papers/FindingCulpritsDySpAN08.pdf
Slides available: .../~sahai/Presentations/FindingCulpritsDySpAN08.pdf
This Handout: .../~sahai/Presentations/FindingCulpritsDySpAN08.H.pdf
Further discussion, related work, and references can be found in the paper.

1 Motivation

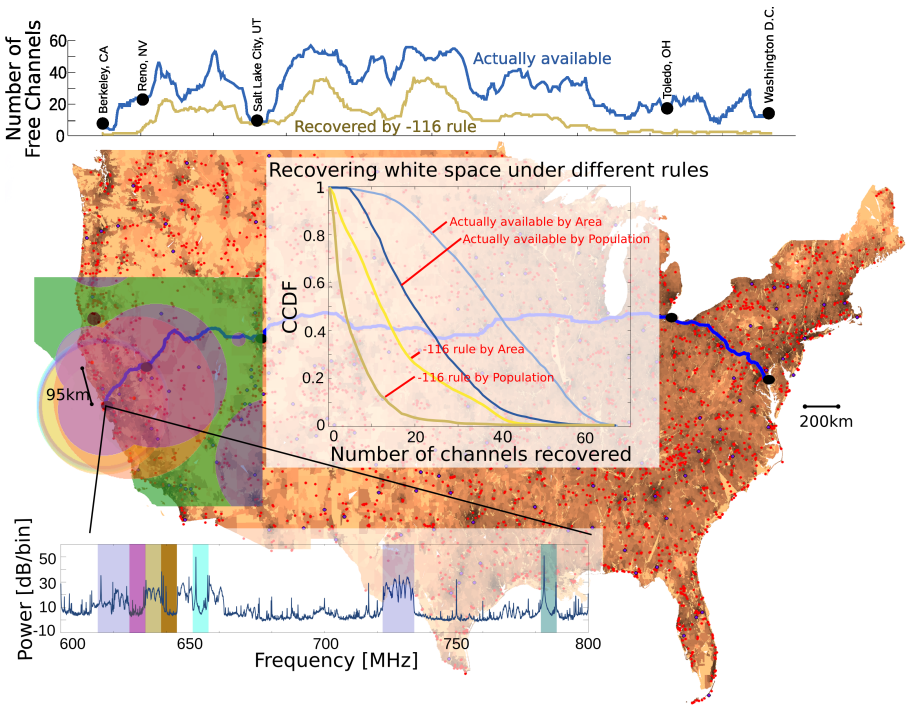

As we take a virtual trip along Interstate 80, we see from the plot at the top that there are lots of channels not being used throughout the U.S. However, the central figure shows that there are somewhat fewer opportunities where most people live. The "sensitivity-based" approach to certifying dynamic spectrum access seems doomed to poor performance. This figure was created by Mubaraq Mishra and is borrowed from A. Sahai, S.M. Mishra, R. Tandra, and K. Woyach, "Cognitive Radios for Spectrum Sharing," to appear in the *DSP Applications* column in the *IEEE Signal Processing Magazine* for January 2009.

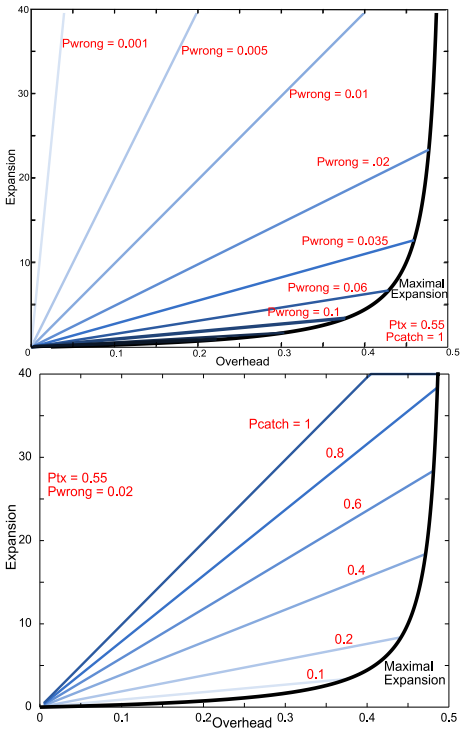

Suppose we catch and send misbehaving users to spectrum jail where *they are banned from using any band for a while, including their own home band*. If the regulator can make errors, even an honest cognitive user will spend some time in jail. We define the punishment overhead as *the average amount of available bandwidth (spectrum the primary is not currently using) the cognitive user cannot recover due to time spent in jail* and the expansion as *the degree to which a cognitive user is willing to attempt to recover spare bandwidth for itself.* Notice that for low overhead with large expansions, the probability *Pwrong* of *wrongful conviction* must be extremely small while the probability *Pcatch* of *getting caught* must be substantial. These figures were created by Kristen Woyach and are borrowed from K. Woyach, A. Sahai, G. Atia, and V. Saligrama, "Crime and Punishment for Cognitive Radios," *Allerton*, Sep. 2008.

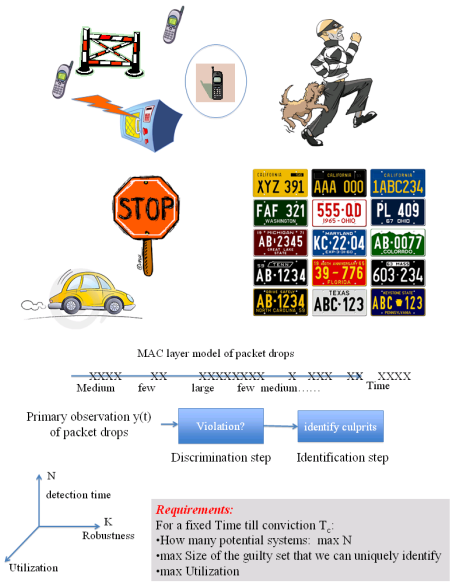

But how can we catch spectrum violators with such confidence? This is the problem of "Hit and Run Radios" identified by Faulhaber. It requires a way to give an identity to each user that can subsequently be detected. Rather than heavy-handedly specifying a specific identity beacon and requiring every primary user to be able to decode it, we propose a "light-handed" approach in which cheating users only need to be identifiable by their pattern of interference. The problem breaks into two parts: discrimination where *we distinguish between natural fading and culpable interference* and identification where *we decide which potential user is actually liable.*

2 Identity System

2.1 Discriminating interference from fading

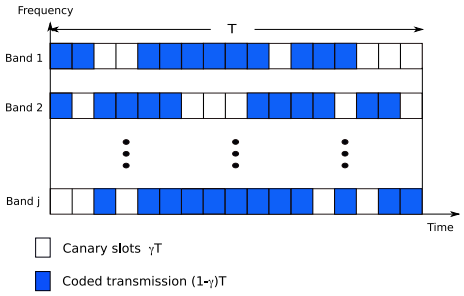

To distinguish between the uncertain background losses and the presence of harmful secondary users, the idea is to introduce canary slots which are *silence slots during which secondary transmission is taboo*. This is checked at device-certification. Canary slots are different in different bands. This allows systems to hop among different frequency bands to maintain stable links. T is *the time till we know that something unnatural is going on* and γ represents *the associated discrimination overhead — the fraction of canary slots required*.

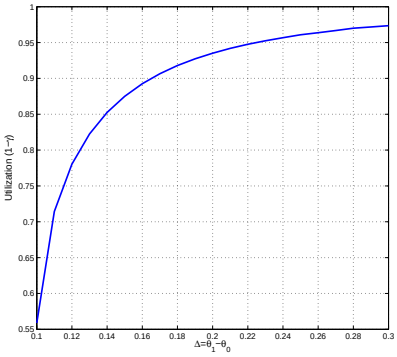

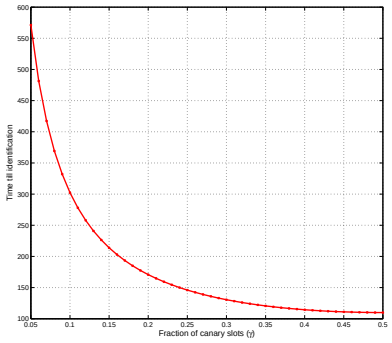

Utilization versus hypothesis separation for catch probability $Pcatch = 90\%$, false alarm probability $P_F = 20\%$, background loss 30% and time till discrimination $T = 450$ slots. Since it is harder to detect subtle interference that slightly degrades primary performance, a larger fraction of silence slots must be used resulting in extra overhead for a fixed detection time. Similarly, extra overhead must be paid if we wish to reduce the time to discrimination.

2.2 Catching Culprits through Interference Fingerprints

If there is only one system then there is a fear of being caught. However when there are many coexisting systems there is an incentive to cheat and hide in the crowd. Hence there is a need for the traceability of culprits.

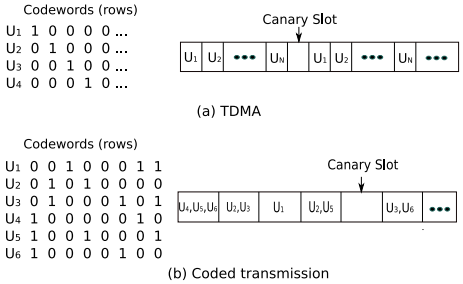

An obvious choice would be to allow for exactly one user to transmit at each time slot (TDMA). Even though this guarantees traceability, it results in very poor utilization. So, why not 'allocate' the same degrees of freedom to many different potential secondary users. This remains easy to certify, but by doing so, higher utilization is achievable. In the given example with $N = 6$ users the average per-user utilization with TDMA is 16.67% in contrast to 35.71% with coded transmission. In each case, we show the code structure and the available transmission slots for the different users.

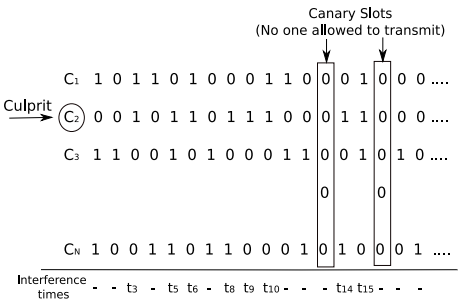

Every cognitive user is assigned a different binary codeword that *defines its allowable transmission slots*. The goal is to identify the guilty parties by matching the interference pattern to the known collection of codewords. The canary slots for discrimination are common to all systems, and beyond that, the i-th user is only allowed to transmit during periods where its code $c_i = 1$. The interference pattern can then be used to identify the culprits as it has the signature of the culprit set, in this case user number 2.

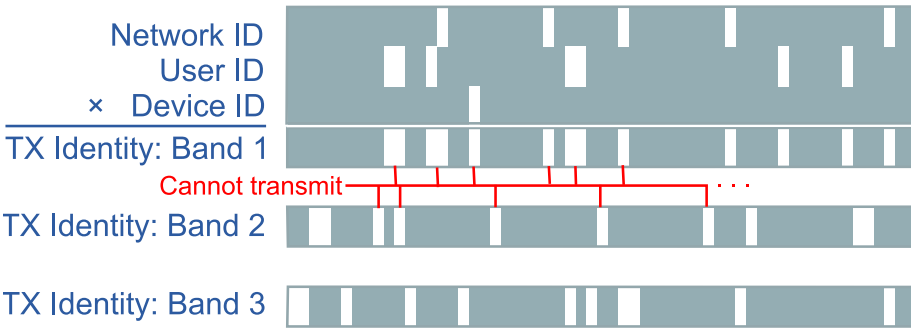

This code can be constructed in a hierarchical manner by ANDing together various codes corresponding to different levels of identity (user, device, network) — if any of the patterns say a slot cannot be used, the overall identity will not use that slot.

Important Variables and a Deterministic Model

- The size of the guilty set K is *the number of simultaneous interferers*.
- Time to Conviction (T_c) measures *how long it takes to uniquely identify the guilty set of culprits*.
- N is the *number of potential users or distinct identities*.
- Secondary Utilization p is the *fraction of time a user is allowed access to a specific band*. Note that there are two kinds of overhead. First are the inverse-canary slots which are common to all users. Second are the silence periods that are different among different users because of their own codes.

Deterministic packet loss model:

- A culprit $g_i = 1$ transmits whenever he is allowed to transmit $c_i = 1$.
- Primary packets are dropped whenever there is an illegal transmission.
- No background losses due to natural causes.

$$y(t) = \bigvee_{i=1}^{N} (g_i \wedge c_i(t)), \ t = 1, \ldots T_c$$

where $\vee$ denotes the (OR) operation and $\wedge$ the (AND) operation. The above equation simply states that a primary packet is dropped if any user from the culprits set is allowed to transmit. This is the simplest possible case.

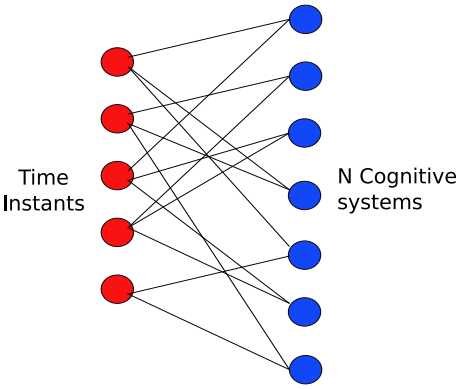

A bipartite graph to model time instants during which different users are allowed transmission. Graphical approaches could then be used to find a set of nodes providing unique coverage of any set of size K, i.e. the number of interferers. This connects this problem to the issue of superimposed codes in graph theory and this connection is used to prove achievable bounds and tradeoffs between the fundamental parameters of interest: T_c, p, N, and K.

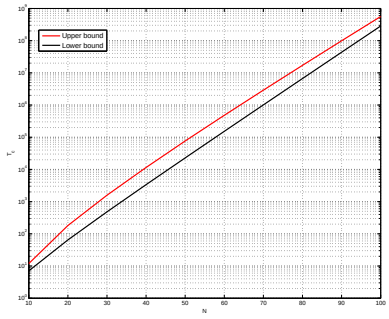

Sufficient and Necessary conditions for time to conviction with $\alpha = \frac{K}{N} = 0.2$ constant plotted versus the total number of potential cognitive systems N. Supporting a large number of potential users becomes prohibitively expensive if we expect a constant fraction of the potential users might simultaneously cheat.

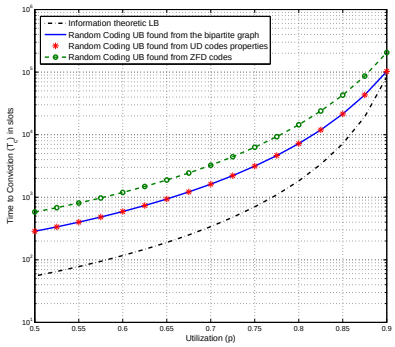

The utilization (p)/time to conviction (T_c) tradeoff with random coding with N=40 and the number of interferers at K=4. This shows a fundamental trade-off between p efficiency, and timeliness for fixed N and K. Intuitively, one could think of a scenario where more slots are added during which all users are allowed to transmit. This leads to an increased utilization but effectively these slots will not help with the conviction process.

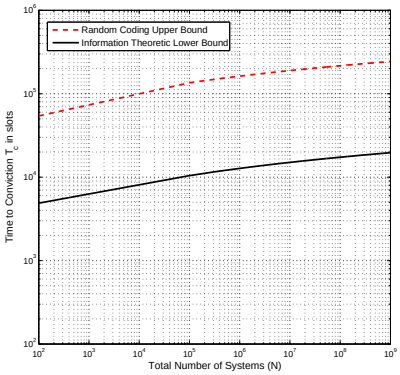

Time till conviction T_c versus the potential number of systems N, with $K = 6$. For a sparse number of culprits, we can increase the potential number of users N and still be able to detect the culprits within detection time constraints.

This talk is intended to bring out the following ideas:

To deter harmful interference in dynamic spectrum sharing, there has to be a way of identifying violators. We propose that this should be done by giving each user their own signature of time slots so that the identity of the interferer is visible directly from the pattern of interference itself. In particular, this helps us distinguish between true interference and random fading. However, this additional capability of being able to catch violators does not come for free. An overhead must be paid in terms of the proportion of slots that we must give up and leave idle. This overhead increases if we want to:

- Increase the sensitivity of detection (How weak of an interferer do we still want to catch)
- Reduce the time to detection (How fast do we want to catch them)
- Increase the number of distinct identities (potential users)
- Increase the resolution in terms of being able to identify multiple simultaneous violators.

Stay tuned: At Asilomar we will reveal how to extend the identification results to a more probabilistic model of loss.