Security in Cognitive Radio Networks: Threats and Mitigation

T. Charles Clancy[1,2] tcc@umd.edu
Nathan Goergen[1] goergen@umd.edu

[1] Electrical and Computer Engineering, University of Maryland, College Park
[2] Laboratory for Telecommunications Sciences, US Department of Defense

***Abstract*— This paper describes a new class of attacks specific to cognitive radio networks. Wireless devices that can *learn* from their environment can also be *taught* things by malicious elements of their environment. By putting artificial intelligence in charge of wireless network devices, we are allowing unanticipated, emergent behavior, fitting a perhaps distorted or manipulated level of optimality. The state space for a cognitive radio is made up of a variety of learned beliefs and current sensor inputs. By manipulating radio sensor inputs, an adversary can affect the beliefs of a radio, and consequently its behavior.**

In this paper we focus primarily on PHY-layer issues, describing several classes of attacks and giving specific examples for dynamic spectrum access and adaptive radio scenarios. These attacks demonstrate the capabilities of an attacker who can manipulate the spectral environment when a radio is learning. The most powerful of which is a self-propagating AI virus that could interactively teach radios to become malicious. We then describe some approaches for mitigating the effectiveness of these attacks by instilling some level of "common sense" into radio systems, and requiring learned beliefs to expire and be relearned. Lastly we provide a road-map for extending these ideas to higher layers in the network stack.

I. Introduction

Cognitive radio offers the promise of intelligent radios that can learn from and adapt to their environment. Much research is currently underway developing various reasoning and learning algorithms that allow cognitive radios to operate optimally in a large variety of different situations.

However, as with many new technologies, initial research has not focused on security aspects of cognitive radio. Typically security is always "bolted on" after the fact by adding some sort of link authentication and encryption. This typically works well for data traversing a wireless network, but not necessarily for things fundamental to the operation of the wireless link itself.

Since cognitive radios can adapt to their environment and change how they communicate, it's crucial that they select optimal, secure means of communications. Data integrity and confidentiality can be handled by higher-layer cryptographic security, so here we focus on attacks fundamental to the cognitive radio itself, and independent of its higher-layer communications techniques.

By putting artificial intelligence (AI) engines in charge of our wireless devices, we need to be aware that these engines can be provided false sensory input by adversaries, and this false input affects its beliefs and behavior (see figure 1). We need to look at threats we would ordinarily see in social networks, rather than computer networks. We define three classes of attacks: *sensory manipulation attacks* against policy radios, *belief manipulation attacks* against learning radios, and self-propagating behavior leading to *cognitive radio viruses*. All types of attacks manipulate the behavior of a cognitive radio system such that it acts either suboptimally or even maliciously.

Protecting against attacks like these cannot be done through cryptographic means. It involves imparting some amount of intuition and common sense into a cognitive radio that allows it to debunk beliefs that don't make sense. In this paper we explore these ideas.

Very little research has examined new threats to cognitive radio due to their intelligent behavior. Some specific work has been conducted looking at attacks in dynamic spectrum access [1], [2], and was broadened to look at a variety of denial of service attacks against policy radios [3]. High-level requirements for using cognitive radio sensing and intelligence to address cross-layer security problems has been examined [4].

In this paper, we focus primarily on the physical (PHY) layer, and provide a general analysis of threats to different types of cognitive radio (including both policy and learning radios). We then present a rough sketch at how such attacks can be mitigated in cognitive radio implementations, and provide a road-map for extending this analysis to the medium access control (MAC) layers and higher. An ideal cognitive engine would provide optimization across all layers, and the threat model fundamentals would apply to those higher layers as well; however, most of the current cognitive radio technologies focus on the PHY, so we use attacks against the cognitive radio PHY as exemplars.

In the rest of the paper we develop the following threats to a cognitive radio network:

1) sensory input statistics can be altered;
2) faulty sensory input statistics can lead to belief manipulation;

This work was funded by the Laboratory for Telecommunication Sciences, US Department of Defense. The opinions expressed in this document represent those of the authors, and should not be considered an official opinion or endorsement by the Department of Defense, or US Federal Government.

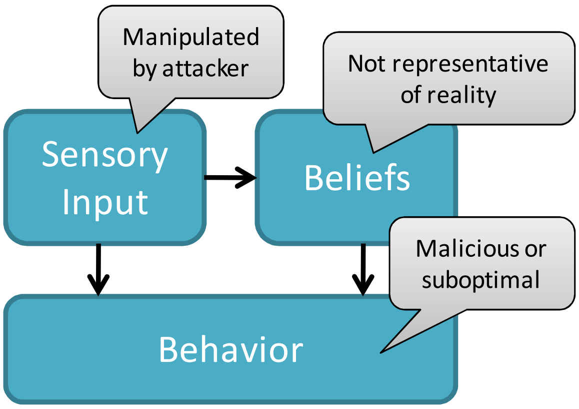

Fig. 1. Relationship between sensor input, beliefs, and behavior in a cognitive engine, showing how an adversary manipulating sensory input can change the beliefs and behavior of a cognitive radio.

3) manipulated individual statistics and beliefs may be distributed through a cognitive radio network; and
4) behavior algorithms based on manipulated statistics and beliefs can result in suboptimal performance or malicious behavior.

To mitigate the effectiveness of these attacks, cognitive radios should:

1) always assume sensory input statistics are “noisy” and subject to manipulation;
2) be programmed with some amount of “common sense” to attempt to validate learned beliefs;
3) compare and validate learned beliefs with other devices on the network;
4) expire learned beliefs to prevent long-term effects of attackers; and
5) attempt to perform learning in known-good environments

Section 2 describes common threat models currently used in security analysis, specifically the Internet threat model, and extensions assumed in wireless networks. Section 3 discusses threats to individual cognitive radio links. Section 4 extends that analysis to networks of cognitive radios. Section 5 details specific attacks against applications of cognitive radio being studied in contemporary literature. Section 6 develops techniques and strategies for mitigating these attacks. Section 7 provides a road-map for further study. Section 8 concludes.

II. Wireless Threat Model

In this section we outline the threat model commonly assumed for wireless networks, to provide a foundation on top of which we can start talking about a cognitive radio threat model.

In any communication networks, there are two major attacker classifications: on-path versus off-path. An off-path adversary can inject data into a stream, or spoof other devices on the network, but cannot, in real time, see the traffic being transmitted. As such, off-path adversaries can be thwarted by using protocols in which devices can only participate if they can see the traffic.

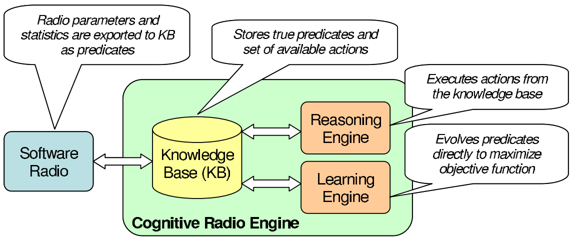
Radio parameters and statistics are exported to KB as predicates
Stores true predicates and set of available actions
Executes actions from the knowledge base
Evolves predicates directly to maximize objective function
Software Radio
Knowledge Base (KB)
Reasoning Engine
Learning Engine
Cognitive Radio Engine

Fig. 2. Components within a single cognitive radio, showing reasoning and learning engines that manipulate the SDR’s operating state

On-path adversaries are by far the most capable. They can both observe and transmit data in real time. This gives them the ability to simply observe traffic, but also spoof, inject, remove, and alter it as well. Protection against denial-of-service (DoS) attacks is difficult, as the adversary can degrade the link such that communication between valid parties is impossible. To protect against non-DoS attacks, a combination of mutual authentication, data integrity protection, and data encryption can be used.

Most security-related protocols design for the worst-case scenario: the on-path adversary. Most transport protocols design for off-path adversaries, and recommend usage of a security sublayer if protection against on-path adversaries is required.

Wireless networks offer additional challenges for protocol designers because it makes being an on-path adversary significantly easier. The link from individual devices to the network infrastructure is exposed much more so than in a traditional wired, switched network. All devices can see all traffic from all other devices within their radio-frequency (RF) range.

Additionally, as in any network, connections between clients and servers typically traverse many physical links, each adding latency. By being so close to one end of the connection, attackers can much more easily spoof packets from devices within the broader infrastructure. Attackers in close proximity to their victim will see packets long before their intended recipient, making it easier to spoof a valid response before the server.

The properties of RF also make deletion and alteration of packets much easier. By simply transmitting a jamming signal, resulting interference will degrade the signal quality for a packet, causing it to be lost. Once deleted, an altered packet can then be spoofed. Also the ease of deletion facilitates denial of service attacks.

With wireless networks, we *must* assume the presence of on-path adversaries. Consequently, most wireless protocols include a built-in link-layer security protocol. For example, IEEE 802.11i [5] provides the necessary mutual authentication, data integrity, and data confidentiality.

As we shall see in the next section, cognitive radio networks further complicate the threat model, as they offer adversaries the ability to interact with devices at an even more fundamental level.

III. Threats to Cognitive Radio

In this section we first describe the components in a generic cognitive radio and explain how they interoperate. We then detail two specific classes of radios: policy radios and learning radios. For these two classes we outline the types of attacks possible to the radios themselves.

A. Cognitive Engine Architecture

A cognitive radio (see Figure 2) consists of four major parts. First is the *software-defined radio* (SDR). This is a highly-configurable wireless communications device, typically capable of synthesizing a large number of communications waveforms by composing processing graphs of different radio components.

For example, an SDR typically has an adjustable front-end allowing it to tune over different frequency ranges and an amplifier that allows communication at many different power levels. The modem components can implement many different modulation types with different symbol rates. Similar flexibility is possible for additional layers, including forward error correction and data framing, multiplexing, and scheduling.

The SDR typically also has a wide variety of sensors. These sensors take digitized RF energy and produce a quantitative result. For example, an energy detector could measure the received power at the tuned frequency, in an effort to determine whether or not that channel is occupied. Specific waveform detectors can do similar tests to determine exactly what type of communications system is occupying the channel. Other sensors could characterize a noise or interference source by measuring the autocorrelation and other useful statistics that could help design an optimal transceiver. Receiver sensors can determine things such as the current signal-to-noise ratio (SNR), bit error rate, and frame error rate.

The SDR then has a programming interface that exposes these configuration options and sensors to a controlling entity. They are essentially a set of inputs and outputs. The controlling entity needs to select the set of inputs that results in optimal outputs, where optimality is often defined with an objective or fitness function. Selecting the radio inputs is then simply a multi-dimensional, discrete optimization problem.

To handle this optimization, a cognitive engine is introduced. All these inputs and outputs show up in the cognitive engine's knowledge base as either read-only (statistics) or read-write (configuration). The knowledge base is a set of logical expressions representing the state of the radio system. Within the cognitive engine, there are two mechanisms for interacting with the knowledge base: the reasoning engine and the learning engine. A *policy radio* only has a reasoning engine, while a *learning radio* has both a reasoning and a learning engine.

The reasoning engine is a set of logical inferencing rules, sometimes called a case-based reasoner. It is provided with a set of actions, the conditions under which those actions are executable, and how those actions affect the state of the knowledge base. The engine then proposes application of these rules in various permutation, searching for a proposed set of actions that will manipulate the knowledge base's state in an optimal way. Here we select a combination of radio configuration values that will maximize some performance metric.

Learning radios are far more flexible, because they include a learning engine. This learning engine is capable of starting with no preprogrammed policy and "trying out" various radio configurations to see how the system performs. For example, a radio can try out different modulation types to see which works optimally in a particular RF environment. Learning radios typically utilize a variety of classic AI learning algorithms, including search algorithms, neural networks, and evolutionary algorithms.

In his dissertation [6], Mitola describes the typical cognition cycle of *Observe* → *Orient* → *Plan* → *Decide* → *Act*. If the radio supports learning, whenever this loop results in a new operating state for the radio, another stage called *Learn* is injected into the cognition cycle that allows the radio to add to its memory information about how the radio transitioned to this new operating state—information that can be used by *Plan* and *Decide* in future cognition cycles. In the context of Mitola's cognition cycle, the goal of our attacks is to manipulate the *Observe* stage, and by doing so we can affect all others. For policy radios, by affecting *Observe* we can influence *Act* for that single cognition cycle. For learning radios, by affecting *Observe* we can influence *Learn* and thereby have a long-term impact on *Act*.

While these AI features allow the radio to be extremely flexible and operate optimally in a large number of different scenarios, they can also offer an avenue for attackers, as described in the following sections.

B. Policy Radios Threats

In a policy radio, the main concern is an attacker spoofing faulty sensor information, causing the radio to select a sub-optimal configuration. Radio sensors take digitized RF and extract useful statistics from it. By manipulating the RF the radio sees, an attacker can cause faulty statistics to appear in the knowledge base.

By understanding how a radio's statistics are calculated, an adversary can manipulate them. Since these statistics operate on raw RF energy, there is no cryptographic means of securing them, as is frequently done to prevent typical communications threats.

We call this term of attacks *sensory manipulation attacks* since they rely on understanding a complex set of logic, and knowing what type of input to provide in order to coerce the desired output.

C. Learning Radios Threats

Learning radios are vulnerable to the same threats as policy radios, where an adversary provides false sensory input. However, a learning radio uses all its experiences to develop long-term behavior, making the attacks much more powerful.

For example, an attacker can introduce a jamming signal whenever a policy radio switches to a faster modulation rate, forcing it to always operate at lower modulation rates, resulting in lower link speeds. This will cause link degradation for

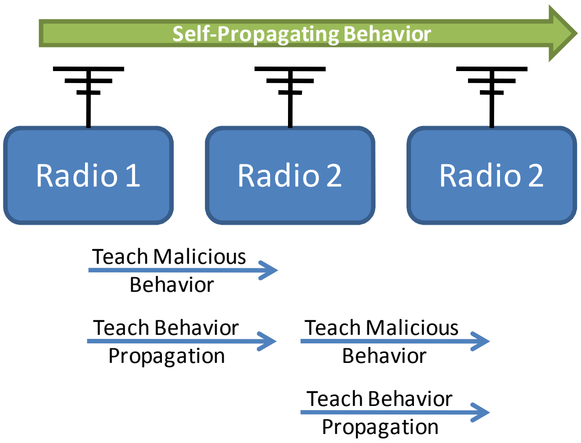
Self-Propagating Behavior
Radio 1
Radio 2
Radio 2
Teach Malicious Behavior
Teach Behavior Propagation
Teach Malicious Behavior
Teach Behavior Propagation

Fig. 3. Process by which malicious, learned behavior, and the ability to teach that behavior to others, can be propagated through a cognitive radio network

the duration of the attack. However a learning radio might permanently associate higher modulation rates with lower data rates, forcing it to always use lower data rates.

Since these attacks can have much longer-term effects on learning radios we term them *belief manipulation attacks*. If you can metaphorically convince a radio that "up is down", and "down is up", and you can seriously impact how it behaves and reacts to particular situations.

In the next sections we describe what effect these attacks can have on a network of cognitive radios, and then describe some technique for mitigating the attacks.

D. *Self-Propagating Behavior*

One of the most powerful types of attacks revolves around the idea of self-propagating behavior. In such an attack, state $\mathcal{S}$ on radio R_0 will cause behavior that induces state $\mathcal{S}$ on radio R_1. Once radio R_1 is in state $\mathcal{S}$, it can induce that same state on radio R_2, and so on. Eventually state $\mathcal{S}$ will propagate through all radios in a particular area.

More generally, it may not just be a single state $\mathcal{S}$, but rather a series of state transitions $\mathcal{S}_1$, ..., $\mathcal{S}_N$ that induce the same pattern of state transitions in neighboring radios.

The result is effectively a *cognitive radio virus* as shown in figure 3, the severity of which will depend on possible side effects of these states. A key feature of these types of attacks is that they can spread between non-cooperative radios that never have direct protocol interaction.

The possibility of such a set of states is not entirely implausible. For example, if many devices are trying to optimally and fairly share a distributed resource, often the optimal behavior is for each device to use an equal fraction of the available resource. When acting optimally, all devices are traversing the same series of states and executing the same behavior. An adversary may be able to influence this equilibrium such that the asymptotic state is not optimal, and possibly even malicious.

IV. Cognitive Radio Networks

From an artificial intelligence perspective, a network of cognitive radio is a set of independent, logical agents, interacting by each observing and changing their environment in some way. These agents can be cooperative, in which case they typically have a control channel for exchanging state information, or non-cooperative, where they act independently.

A simple two-node cognitive radio network consisting of a single, point-to-point, duplex link is a simple example of a cooperative network. The two radios each configure themselves to optimally communicate, and share necessary channel statistics in order to jointly come to the optimal configuration. The transmitter is responsible for the configuration and must communicate this to the receiver, and the receiver measures the channel statistics and must communicate these to the transmitter.

In larger networks, each node has its own sensors and gathers its own statistics. Each pairwise communications link must have a common configuration in order for data transfer to be possible. A simple example of this is the proposed IEEE 802.22 standard [7], where each client devices makes channel measurements and sends them to the base station. The base station then makes the determination on which channel is optimal and least likely to be occupied by a legacy communications system such as a television station.

In a non-cooperative network, an attack against one cognitive radio will not affect any others, because other devices will independently take their own sensor measurements and make their own decisions.

In a cooperative network, attacks against a subset of nodes can have further-reaching effects. For example, a conservative implementation of IEEE 802.22 would have logic that causes all devices to migrate to a new frequency if a single device detected a television signal. Thus an adversary can spoof a television signal to a single IEEE 802.22 device, and the network will migrate everyone to a new frequency, allowing the attacker contention-free access to the spectrum.

More security-conscious implementations will try to guard against byzantine failures by more intelligently fusing conflicting statistics from a distributed sensor network. Some strategies for doing this are discussed in later sections.

V. Classes of Attacks

In this section we describe specific scenarios an attacker can construct by manipulating knowledge base state on devices in a cognitive radio network. We relate these attacks to some of the common applications of cognitive radio currently under study today.

A. *Dynamic Spectrum Access Attacks*

The first type of attack, which was introduced in [1], [2], is called the *Primary User Emulation* (PUE) attack, and can be effective in dynamic spectrum access (DSA) environments. In such environments, a primary user owns a license to a particular frequency band, and can use it whenever they wish. When they are idle, secondary devices can opportunistically use the available spectrum. Such secondary devices need spectrum sensing algorithms to detect when the primary user is active.

All an attacker need do is create a waveform sufficiently similar to that of the primary user to trigger a false positive in the spectrum sensing algorithm. The secondary devices within range will believe a primary user is active, and will cause the system to vacate the channel. This gives the adversary unrivaled access to the frequency band.

Fortunately the effects of this attack are transient, as it is only a sensory-manipulation attack. Once the attacker vacates the frequency, the secondary users notice the spectrum being once again idle, and can resume using it.

Other DSA algorithms are more stateful, and accrue more detailed statistics about primary users. For example, some DSA algorithms gather channel access statistics for primary users in an attempt to predict when the channel will be idle, based on current *and past* behavior of the primary user [8]. Thus spoofing primary user waveforms can affect the long-term behavior of a secondary user, turning this attack into a belief-manipulation attack.

If an adversary wishes to deny service to a secondary user operating in the presence of a time-division multiple access (TDMA) primary user, the attacker needs to make the primary user's access pattern look random during the learning phase of the secondary user, rather than periodic. As a result, a secondary user cannot derive very much information about when they can transmit without interfering, significantly decreasing their capacity if not completely preventing transmission. This memory will persist after the attacker discontinues transmission.

B. Objective Function Attacks

In adaptive radio, the cognitive engine has a large number of radio parameters under its control. The cognitive engine manipulates these parameters over time in an effort to maximize its multi-goal objective functions. These attacks apply to any learning algorithms that utilize objective functions, most notably various forms of hill climbing and genetic algorithms. Since these attacks apply to learning engines, they are *belief-manipulation attacks*.

Some possible input parameters could be center frequency, bandwidth, power, modulation type, coding rate, channel access protocol, encryption type, and frame size. The radio might then have three goals: *low-power*, *high-rate*, and *secure* communication. Depending on the application, each of these goals has a different weight. For example, if using the system for instant messages or email, *low-power* and *secure* would have higher weights than *high-rate*. For voice or video applications, *high-rate* and *secure* would have higher weights than secure.

There are a few different types of attacks possible on a system like this. When radios are in a learning phase, they try different combinations of input parameters, measure the observed statistics such as bit error rate, and then evaluate the objective functions to see which inputs give the best results for their application.

Of the three goals – *low-power*, *high-rate*, and *secure* – only one is affected by the channel. *Low-power* and *secure* are defined directly by inputs, while *high-rate* is defined by system outputs. Thus by affecting the channel, and adversary can manipulate whether or not high-rate communications is achieved.

More concretely, imagine the following objective function

$$f = w_1 P + w_2 R + w_3 S \tag{1}$$

where w_i are the weights and P, R, and S represent the three goals of power, rate, and security. Imagine an adversary wishes to force a radio to use some security level s_1 rather than the more secure version s_2, where $s_1 < s_2$. Whenever the cognitive engine tries using s_2, the adversary can jam the channel, artificially decreasing R from r_2 to r_1 with $r_1 < r_2$. In particular, an adversary would need to cause sufficient interference such that

$$w_1 P + w_2 r_2 + w_3 s_1 > w_1 P + w_2 r_1 + w_3 s_2 \tag{2}$$

or solving for r_1:

$$r_1 < r_2 - \frac{w_3}{w_2}(s_2 - s_1) \tag{3}$$

The consequence of such an attack is that whenever a higher security level is attempted, the system's objective function decreases, and that higher security level is never used.

This type of attack can already be successfully executed on static systems that involve human configuration. For example, consider an engineer trying to configure a point-to-point wireless link to interconnect two networks. If every time the network is activated with encryption enabled, an adversary jams the network, the engineer may eventually give up and believe there is a crypto-related impediment to using encryption, and simply run the network without encryption.

Similarly, the IEEE 802.11u standards are defining weaker security mechanisms for wireless local area networks (WLANs) to support placing E911 phone calls via unauthenticated WLANs in emergency situations. This represents a case where the need for connectivity outweighs security, and the logic behind this could be exploited if implemented with a learning engine.

Additionally, similar attacks can be used to cause radios to believe that certain frequencies, bandwidths, or modulation types are less optimal and should be avoided. This would allow an attacker to sculpt the waveforms used by a cognitive radio to suit its goals, whatever they may be.

It is very important to note that this type of attack only works when the radio is performing *online* learning, that is the radio is performing some sort of online optimization of the search space. Radios that perform *offline* learning observe the environment once, and then perform an offline search to find the optimal configuration; such radios assume independence of their observations and configuration, and consequently are immune to these attacks. It should also be noted that such radios theoretically do not require a learning engine, and their behavior can be reduced to the case-based reasoning of a policy radio [9].

C. Malicious Behavior Attacks

In this section we discuss an extension of the objective function attacks, where we teach a radio to become unknowingly malicious. Here we examine a scenario that cause a cognitive radio to become a jammer.

Consider a system where a primary user is intermittently accessing a channel. Secondary users have channel sensing algorithms that can detect primary and secondary users access the channel. They have an objective function that balances throughput T and interference I, and looks like

$$f = w_1 T - w_2 I \tag{4}$$

The system seeks to maximize throughput while minimizing interference.

The desired result is a secondary user will only communicate when the primary user is idle. However, if an adversary uses a jamming waveform that cannot be detected by the secondary user's sensing algorithms, he can artificially decrease T when the primary user is idle. As a result, the cognitive radio will learn that the only time it can achieve useful communications is when the primary user is active. This effectively turns the cognitive radio into a jammer.

Accomplishing this should be fairly straight-forward using common commercial waveforms. Imagine a primary user uses some linear, narrow-band modulation scheme, and the secondary radio uses Orthogonal Frequency Division Multiplexing (OFDM). Typically OFDM uses pilot tones on a few of the subcarriers for receiver channel estimation and synchronization. By transmitting a carrier wave on those pilot tones, an adversary can prevent OFDM receiver synchronization and consequently block any useful communication. Additionally, these CW signals are unlikely to trigger the cognitive radio's spectrum sensing algorithms, which would prevent detection of the attack.

VI. Attack Mitigation

In this section we describe a variety of techniques for mitigating the effectiveness of the attacks detailed in the previous sections.

A. Robust Sensory Input

Improving sensor input can significantly help reduce the gullibility of cognitive radios. For example, if radios could carefully characterize the difference between interference and noise, they could distinguish between natural and man-made RF events. Such sensors could also feed specialized policy engine subroutines that specifically look for hostile signals that may be attempting to corrupt a radio's beliefs.

In a distributed environment, a network of cognitive radios can fuse sensor data to improve performance. For example if multiple cognitive radios exchanged time-synchronized, digitized RF, they could run cross-correlation algorithms to more precisely determine the difference between an attacker and noise.

All sensory input should be considered "noisy", since even without the presence of an attacker, statistics can occasionally be incorrect. For each input, cognitive radio designers should quantify the probability of detection failure in both benign and hostile environments. In some scenarios, attackers may be power-limited, allowing designers to compute theoretical upper-bounds to their effectiveness.

B. Mitigation in Individual Radios

In order to mitigate attacks against individual radios, we need to instill some amount of "common sense" into radio systems.

Protecting individual policy radios from attack is difficult. They look at the current environment and evaluate what course of action to pursue subject to their policy. An adversary with knowledge of the policy could provide false sensor inputs in an effort to affect the radio's behavior. Even without knowledge of the policy, an adversary could use various fuzzing techniques to infer the policy.

As a result, radio policies should be carefully evaluated to protect against malicious sensor input. All possible states can be enumerated as the product space of all possible knowledge base values. The radio policy then defines a state transition framework that can be overlaid on the state space. Formal state-space validation, as is often done with cryptographic network protocols, can then be applied to the resulting state machine to ensure that a "bad state" is never reached.

However, in scenarios such as the PUE attack, no carefully-engineered policy can protect a simple radio from detecting an adversary as a primary user. Mitigating such attacks relies on developing better sensing algorithms with lower false-positive rates. A better-developed sensing algorithm may be able to distinguish the legitimate primary user from the adversary. Additionally, radios might compare their perception of their environment with characteristics know *a priori*. For example, in [2], the authors attempt to geolocate primary transmitters, and compare their measured location to the locations of known primary users.

Protecting individual learning radios is also difficult, but there are some strategies for preventing attackers from indefinitely altering a radio's beliefs.

First, beliefs must always be under constant reevaluation. A feedback loop should constantly be updating learned relationships between cognitive radio inputs and outputs. For example, coefficients in a neural network should be constantly recomputed, using a moving average, such that previously-learned behavior that may be incorrect can be expired.

If continuously-learning radios are impractical or not desired for a particular scenario, then the learning phases *must* be conducted in a truly controlled environment. Outside auditing must take place to ensure that no adversarial signals are present during the learning phase.

Another approach would be to build in logic that could invalidate learned actions that were known to violate certain principles. Writing such rules is difficult, because learned actions can often be non-intuitive (hence why using artificial intelligence can help us solve tough problems). For example, a radio might discover that achieved capacity is inversely proportional to bandwidth, which contradicts Shannon's theorems. However, this might be caused by inadvertent overlapping with an adjacent-channel interferer.

C. Mitigation in Networks

In a network of cognitive radios, where we assume there is some sort of control-channel connectivity between cognitive engines, mitigating possible threats becomes a very interesting space. We now have a group of independent AI agents each seeking to maximize their own performance, and possibly global performance, depending on their level of cooperation with the devices around them.

This new-found connectivity at the cognitive-engine level between peers allows us to use various techniques from *swarm intelligence* [10]. Swarm intelligence is a set of algorithms that mimic various animal behaviors; a common example is ant colony optimization that seeks to find a goal by sending virtual ants off searching for the goal, and leaving pheromone trails back to their starting point after finding intermediate solutions.

One technique of specific applicability to security is particle swarm optimization (PSO) [11]. Each cognitive radio in a network represents a particle, each with it's own hypothesis about what the best behavior is in a particular situation. The behavior it selects, however, is not wholly dependent on its own hypothesis, but is actually a weighted average of all hypotheses in the network.

An example application of PSO is to the PUE attack against DSA systems. Here each device could make its own hypothesis about whether a particular transmission was from a primary user, and the group majority would then be the group decision. Certainly a weighted majority could be taken, based on which devices are most likely to have the most accurate hypothesis (due to proximity, sensor capabilities, etc).

This approach could also be used in adaptive radio. For example, in the scenarios we described earlier, radios used various learning algorithms to determine how their inputs affected their objective function. An attacker could manipulate that process to cause altered behavior in the radio. If PSO were used during this learning phase, radios would experience something more akin to group learning, rather than individual learning, making it more difficult for an attacker to influence the system.

Note, however, that such approaches should be used with care. By allowing this group learning and decision making, an attacker with influence over a few devices may now be able to affect the outcome of the entire group.

VII. Road-map for Future Research

In this paper we primarily focused on PHY security. We discussed scenarios where cognitive radios had control over and sensory input from PHY-level components, such as energy detectors, signal classifiers, modem components, and error correcting codes. However, there's no reason this cannot be extended to the MAC, routing, transport, or even application layers. The eventual goal of cognitive radio is to permeate the entire network stack to allow intelligent control and security to an entire system [4]. The emphasis on the PHY layer so far is mostly because that's as far as basic research into cognitive radio has been formally defined.

A. Understanding Identity

The next concept necessary to extend cognitive radios' security capabilities is to allow the cognitive engine to reason about the identity of things in its environment. Currently cognitive engines may understand that energy at some frequency is another wireless device, and may even be able to tell the difference between different classes of devices, such as primary and secondary users, but the cognitive engine itself cannot recognize and differentiate two different secondary users.

Adding this recognition can be accomplished in two different ways. First, more PHY-layer techniques could be used, such as RF fingerprinting. This would extend the current sensor capabilities to more fine-grained differentiation between transmitters, and give them the ability to recognize transmitter's they've seen before. No detection rules are perfect, however, and this approach may still allow an attacker to manipulate their RF fingerprint to be detected as someone or something else.

A better approach is to combine PHY-layer approaches with MAC-layer processing. For devices within the cognitive radio's network, the MAC layer can provide much more specific and reliable information, especially if MAC-layer protocols are cryptographically protected to prevent manipulation. By feeding this information to the cognitive engine, a radio can now reason about its neighbors, because it knows who they are and when they're communicating. This allows it to build up learned statistics specific to each neighbor, rather than for all neighbors.

B. Earning and Using Trust

Once a cognitive engine can reason about the identity of its neighbors, it can start learning trust metrics, which would in turn allow radios to reason about the trustworthiness of their neighbors. Trust is a loaded word that means many different things in many different applications. At the MAC layer, trust typically involves some quantification of how fairly a device is accessing the shared medium. If a device is greedily using more than its fair share, or is maliciously trying to prevent other devices from communicating, it could be seen as an adversarial force in the network. Also its willingness to receive packets is important.

In a MAC layer, cognitive radios use feedback from the PHY and MAC layers to quantify how fair and reliable different neighbors are. They can then use this information to who can communicate to whom, and when (i.e. the essence of medium access control).

One example scenario is a new device wishing to join a cooperative cognitive radio network. Initially the devices in this network have no knowledge of this new device or reason to communicate with it. If, over time, this device proves itself to be trustworthy, then other nodes in the network will grant it more access to the shared spectrum.

This metric of MAC trust also could be used for distributed rate control. More trusted devices can be granted more capacity than they are entitled to for transient situations, however if this is abused their trust rating with neighbors may decrease, making it less likely that they'll be given extra capacity in the future.

C. Trust in Networking and Routing

The next logical step is to extend reasoning about the trustworthiness of neighbors in a MAC layer to reasoning about devices further away in the network. To accomplish this, a cognitive engine needs control over and feedback from a system's network layer. For example, the cognitive engine needs statistics from each transmitted and received packet, in addition to the ability to affect decisions about L3 communications and routing.

This would allow a cognitive radio to learn trust information about the broader network and use it to develop policy about routing and forwarding of traffic. Certainly trust metrics in ad hoc routing is far from a new concept [12], [13], but allowing a cognitive engine to control how that trust affects the system could allow for new, emergent behavior implementing policies never before considered.

As cognitive radio propagates up the network stack, new challenges will arise. In particular, adversaries will have more "sensors" they can attempt to fool. However one advantage to MAC-layer and higher sensors is that in many cases the data they are examining can be cryptographically protected to prevent adversarial manipulation.

VIII. Conclusion

In this paper we examined the problems associated with adversarial manipulation of cognitive radio sensory inputs, in an effort to cause the victim cognitive radio to behave suboptimally or maliciously. We showed how this ability to manipulate sensors introduces new threats beyond those typically associated with a wireless network.

Some of the most powerful attacks involve an adversary manipulating the RF spectrum while a victim cognitive radio is in its learning phase. By skewing the radio's performance in certain situations, the attacker can make a victim radio believe things that are untrue, and consequently cause it to behave in a suboptimal or malicious manner. Any radio that can *learn* from its environment can also be *taught* by its environment. We demonstrated several different attacks against dynamic spectrum access and agile radios, and even described how self-propagating behavior could lead to a *cognitive radio virus*.

Mitigating such attacks involves building checks and balances into a cognitive radio. First, learned beliefs should never be permanent. Otherwise it could indefinitely induce undesired behavior. Second, cognitive radios should always assume sensors are noisy, whether its due to nature or a malicious force. Lastly, cognitive radios could collaborate with other cooperative radios on the network to coherently develop beliefs that are less likely to be subject to manipulation.

Most of the attacks presented are specific to PHY-layer cognition. Sensors and decisions at higher layers can often be cryptographically protected, making them less vulnerable to manipulatory attacks. We lay out a road-map for how a cognitive engine could help provide additional security at the higher layers, focusing primarily on extending cognition to provide security at the MAC layer.

We're a long way away from Mitola's famous example where a military radio can realize it's been lost and been picked up by an allied soldier, resulting in the deletion of its keying material and immediately connecting the ally to the appropriate liaison office [4]. However through incremental advancements we can build cognition capabilities into the entire network stack. Security is an important part of that – both recognizing threats to cognitive radios, but also how cognition can improve security – and this paper provides the initial building blocks for such advancement.

References

[1] R. Chen and J. Park, "Ensuring trustworthy spectrum sensing in cognitive radio networks," IEEE Workshop on Networking Technologies for SDR 2006.

[2] R. Chen, J. Park, and J. Reed, "Defense against primary user emulation attacks in cognitive radio networks," *IEEE Journal on Selected Areas in Communications*, 2007.

[3] T. Brown and A. Sethi, "Potential cognitive radio denial-of-service vulnerabilities and protection countermeasures: A multi-dimensional analysis and assessment," International Conference on Cognitive Radio Oriented Wireess Networks and Communications, CrownCom 2008.

[4] J. Mitola, "Cognitive INFOSEC," IEEE Wireless and Microwave Technology Symposium 2003.

[5] IEEE 802.11i, "Wlan specifications for enhanced security." 2004.

[6] J. Mitola, "Cognitive radio: An integrated agent architecture for software defined radio." Ph.D. Dissertation, KTH, 2000.

[7] IEEE 802.22, "Working group on wireless regional area networks." http://www.ieee802.org/22/.

[8] T. Clancy and B. Walker, "Predictive dynamic spectrum access," SDR Forum Technical Conference 2006.

[9] T. Clancy, J. Hecker, E. Stuntebeck, and T. O'Shea, "Applications of machine learning to cognitive radio networks," *IEEE Wireless Communications*, August 2007.

[10] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natrual to ARtificial Systems*. Oxford University Press, 1999.

[11] J. Kennedy and R. Eberhart, *Swarm Intelligence*. Morgan Kaufmann, 2001.

[12] C. Carlof and D. Wagner, "Secure routing in wireless sensor networks: Attacks and countermeasures," IEEE Workshop on Sensor Network Protocols and Applications.

[13] Y. Hu, D. Johnson, and A. Perrig, "Sead: Secure efficient distance vector routing for mobile wireless ad hoc networks," IEEE Workshop on Mobile Computing Systems and Applications 2002.