Potential Cognitive Radio Denial-of-Service Vulnerabilities and Protection Countermeasures: A Multi-dimensional Analysis and Assessment

Timothy X Brown, Amita Sethi
Interdisciplinary Telecommunications
Electrical and Computer Engineering
University of Colorado, Boulder, CO 80309
timxb, sethi @colorado.edu

Abstract— **Cognitive radios sense spectrum activity and apply spectrum policies in order to make decisions on when and in what bands they may communicate. These activities go beyond what is done when traditional radios communicate. This paper examines the denial of service vulnerabilities that are opened by these additional activities and explores potential protection remedies that can be applied. An analysis of how vulnerable are victim cognitive radios to potential denial of service attacks is presented along different axis, namely the network architecture employed, the spectrum access technique used and the spectrum awareness model. The goal is to assist cognitive radio designers to incorporate effective security measures now in the early stages of cognitive radio development.**

Keywords-cognitive radio; denial of service; vulnerability; countermeasure

I. INTRODUCTION

A cognitive radio (CR) employs software to measure unused portions of the existing wireless spectrum (so-called white space) and adapts the radio's operating characteristics to operate in these unused portions in a manner that limits interference with other devices [8]. Spectrum regulators such as the Federal Communications Commission (FCC) in the United States (US), recognize that CRs can be applied to dynamically reuse white spaces in licensed spectrum bands, thereby efficiently utilizing under-utilized spectrum [FCC02]. A number of research efforts such as the Defense Advanced Research Projects Agency (DARPA) Next Generation (XG) project in the United States [2, 3] and the End-to-End Reconfigurability (E2R) program in Europe [4] are working towards devising techniques for realizing different aspects of cognitive radio devices. The technological advances in CRs are of such a magnitude that the FCC is of the view that none of the other advances "holds greater potential for literally transforming the use of spectrum in the years to come than the development of software-defined and cognitive or "smart" radios" [6].

However, cognitive radios may be susceptible to actions which prevent them from being able to communicate effectively, so-called denial of service (DoS) attacks. These actions might also induce an otherwise legitimate cognitive radio to interfere with a licensed transmitter. In this paper we do not identify the motives for such actions. They could be from one or more malicious agents that wish to prevent a CR from communicating. It could be from a valid CR that is malfunctioning or one that is misconfigured. Whether they are due to malicious, malfunctioning, or misconfigured behavior, the actions are treated equally. Actions such as direct jamming of the CR communication would affect any radio and so are not of interest here. What we seek to understand are the attack vulnerabilities that are enabled *because* of the CR functionality. We further seek to understand to what extent these attacks are more effective than direct jamming of the radio signal. This paper is based on an earlier paper [23] but extends the analysis to assess how the vulnerability of a victim CR varies as a function of CR network architecture, the spectrum access technique and how the CR becomes aware of spectrum usage and availability in its vicinity.

II. TRADITIONAL VERSUS COGNITIVE RADIOS

While a traditional radio allows minimal user interaction and has unalterable receiver transmitter operations, the CR houses advanced functionalities of [8]:

- *remote reconfigurability*: "the capability of adjusting operating parameters for the transmission on the fly without any modification of the hardware components"
- *spectrum sensing*: a CR device senses its radio environment and adapts its mode of operation in response,
- *spectrum policy based operation*: CR devices' spectrum access behavior are confined by a set of policy rules, and
- *geo-location*: the CR determines its geographical coordinates via methods such as GPS.

These additional capabilities enable a CR to identify fallow unlicensed bands and facilitate opportunistic secondary use in these bands without causing interference to primary users.

Both traditional and cognitive radios can work in licensed or unlicensed spectrum bands. However CRs can work in additional bands that include bands that require sensing or subject devices to other band-specific restrictions that can be satisfied by CRs. Additionally, CRs can operate in several licensed models. Apart from operating in unlicensed bands that are free for use by any radio, CRs can operate in licensed bands that allow unlicensed secondary use by any CR under some policy-defined rules and limitations. They may also operate in yet other licensed bands that allow secondary CR use, but only to a specific licensed set of users under specific conditions. Thus

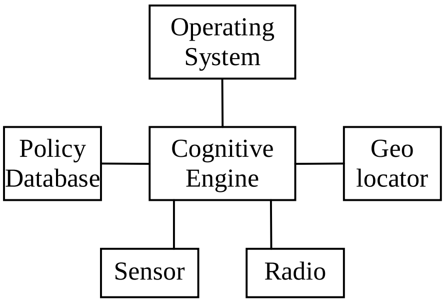

Figure 1: Cognitive radio components.

from a radio user's perspective, a CR has potentially many more opportunities to communicate than traditional radios, and from a spectrum managers' perspective, the CR enables more flexible and targeted spectrum policies.

III. Operational Cognitive Radio Aspects

A. Components

Figure 1 shows the basic components of the cognitive radio. The operating system represents the higher-layer communication functionalities above the radio Physical and Link layers. This generates and receives the traffic information which is to be sent and received by the operating system. A sensing element measures information about the radio environment and provides the information to the cognitive engine. The cognitive engine combines sensor information with policy information to make decisions about when and how it will communicate using the radio transmitter and receiver. Some CRs also depend on knowledge of the transmitter location which is provided by a geolocator such as a GPS receiver.

B. Architectures

CRs can be broadly classified into one of three network architectures as shown in the figure below. They can range from architectures that encompass all six components in a single non-cooperating device to networked architectures where none of the CR components may be co-located with each other. This model includes multiple instances of each component. For example there may be dedicated sensing nodes that communicate with a centralized cognitive engine that then directs remote transmitters on how they can communicate. Furthermore several distributed CRs may choose to share information such as measurements, location, or policy in order to make more informed and coordinated communication decisions. To the cognitive engine, the other CRs are effectively sensing, geolocation, or communication extensions. Many cooperative schemes (centralized or distributed) envision a common control channel that is a well known link to share information [13, 19]. The operational advantages or disadvantages for any architecture in this range are not considered. Rather, we consider the security vulnerabilities when any of these components may or may not be collocated with each other.

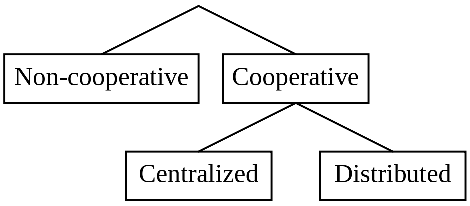

Figure 2: Cognitive radio network architectures.

C. Access Methods

The cognitive radio can operate as an *overlay* or an *underlay*. In an overlay the CR searches for white space bands with which to communicate and generally avoids transmitting power in occupied licensed channels. In an underlay the CR uses spread spectrum or ultra-wideband techniques along with careful power control so as to ensure that no licensed band receives a strong enough signal to cause interference. In addition to communication, a CR may employ these access methods for exchanging policy, location information and sensing measurements as well.

D. Spectrum Awareness

A CR may become aware of spectrum usage and availability in its vicinity through different models. The models considered in this paper are: the Geo-location/database approach, the beacon/control signal approach and the detection/sensing approach [7]. A CR may use the

- *geo-location/database* approach to identify vacant channels at its location, to download applicable policy certificates or to access real-time data of primary users active in its current location.
- *beacon/control signal* approach to announce local policies or a list of spectrum bands that are vacant within the service area of the licensed transmitter.
- *detection/sensing* approach to identify primary spectrum users, to locate a spectrum hole or to detect signals from cooperative radios in the transmission range of the device.

With this overview of elements of CR and the three design axes we turn to the potential denial-of-service vulnerabilities that are possible.

IV. Denial-of-Service Vulnerabilities

A denial-of-service (DoS) attack is an act of preventing authorized access to a system resource or the delaying of system operations and functions [17]. In this paper, it is a denial of communication to legitimate users—the CRs—even when the system resources—such as unused frequencies—are available. Another DoS attack relevant to CRs is when a CR is induced to communicate so that it causes interference with a licensed transmitter. This attack, is also a form of DoS if it leads to a, perceived failure of CR that forestalls the widespread deployment of CR technology, preventing the anticipated benefits to spectrum management from being realized.

A. The Traditional Jamming Attack

A simple denial of service of attack is for an attacker to transmit a continuous high-power signal that prevents usable reception. This brute-force approach can be applied to any type of radio transmission. In general there are approaches such as spread spectrum that can make a radio more robust to these kinds of attacks. The greater the spreading of the signal, the harder it is to detect and the more robust it is to jamming attacks.[1] A cognitive radio has a disadvantage and related advantages to this brute-force jamming attack. Because the cognitive radio is operating in spectrum as available the signal bandwidth may be constrained limiting the protective spreading that is possible. However, the cognitive radio is designed to operate in many different bands. Further it is assumed that the CR generally has robust mechanisms for choosing which band to communicate in the presence of licensed and other users. With these capabilities, an attacker would need to simultaneously jam many different communication bands or have reliable techniques for detecting the CR as it switches between the many bands. A potential complication for the attacker is the presence of the licensed users. The attacker may need to avoid these users as it attempts to detect and attack a target CR. We are especially interested in attacks that use transmission that is not per se prohibited. For instance, an attacker may be able to transmit otherwise legitimate packets that prevent CR communication. Such attackers may seek to disrupt CR communication while operating within legal bounds. Or, the so called attacker may be another legitimate CR whose operation is not compatible with the CR in question. In either case, we seek to understand the vulnerabilities of such attacks.

A direct attack on the signal can be effective. However it makes it easy for the attacker to be detected and countermeasures taken. Therefore an attacker will seek techniques to limit its exposure to countermeasures by reducing the fraction of time and power needed to prevent communication. A concept to capture this notion is *jamming gain* [21]. The attacker has greater jamming gain as it reduces the time or power that it needs to transmit in order to achieve the same effect as with direct jamming. For this paper we will only discuss this in general terms. However, we should be clear that the attacker wants to maximize the jamming gain. The question then is how the different elements of a CR—its architecture how it networks with other CRs; the spectrum access technique it uses and the spectrum awareness model—open up jamming gains to the attacker.

B. *Traditional vs. CR Avenues of Attack*

Traditional jamming occurs at the communication receiver. An attacker which is close to the receiver can jam the communication; potentially with less power than transmitted by the transmitter. An attacker close to the transmitter has no special advantage when jamming the receiver. Knowing when a transmitter is on can be useful information in jamming [21]. But all attacker locations which can sense the transmission are equally effective. With a cognitive radio, the transmitter may also need to receive beacons, location, policy or sensor information and so the attacker can prevent the receiver from receiving by jamming the transmitter Put another way, an attacker near a traditional radio can effectively only interfere with reception. An attacker near a CR can interfere with reception or prevent transmission. Moreover, if the detection/sensing approach is used, it is possible to spoof sensitive detection functions with weak jamming signals. A single spoofing attacker could affect CRs distributed over a larger geographical area. Thus, in comparison to a traditional radio, a CR has greater exposure—it allows more attacks from more places.

C. *Threat Model for CRs*

An attacker is one or more radios that can be in the vicinity of legitimate CR. They can demodulate legitimate signals but can not necessarily decode encrypted messages. Significantly, we assume that the attacker must communicate using otherwise legitimate signals. While this assumption is somewhat restrictive, its main purpose is to avoid considering attacks that simply blanket all potential communication with high power noise; or that cause wanton interference with primary users. As will be seen even with this restriction a large number of potential vulnerabilities arise. The attacker can create different types of signals including the following:

- False signals that can be perceived as primary users' signals
- Messages that can be received by the victim CRs. The messages are not necessarily considered from a legitimate CR user if authentication is used.
- Jamming signal that can prevent messages from being received by a receiver.

The power needed to attack depends on the type of attack. The least power is required to create a spoofing signal that only needs to be detected. More power is needed to create false messages that are correctly received by a victim CR. The highest power is needed for outright jamming that overwhelms other received signals. These powers decrease as the attacker becomes closer. As they decrease the attacker's energy use can be reduced and simpler, lower cost, and smaller RF front ends may be used. Antennas can go from large high gain dishes to compact integrated antennas. As a result, the attacker becomes harder to detect at shorter distances. It is conceivable that the attacking radio might become attached to the victim radio near its antenna from which detection by anyone but the victim itself would be very difficult.

Beyond these radio-based attacks, the attacker may gain access to the victim device's interface and be able to deliberately misconfigure the device or gain access to security passwords. In the worst case they can compromise the node so that it is a malicious participant in the CR communication. These non-radio-based attacks are not part of the threat model in this paper.

V. POTENTIAL CR DOS VULNERABILITIES AND PROTECTION COUNTERMEASURES

We categorize the CR DoS attacks into denial and induce attacks. *Denial* attacks can prevent communication through placing the victim CR in one or more of the following states:

1. All available spectrum appears to be occupied by licensed transmitters
2. No policy is available that enables it to transmit.
3. Location information is unavailable or has too low accuracy.

[1] But, an attacker that is close enough or has a powerful enough transmitter can always detect or attack a spread spectrum signal.

4. The sensor is unavailable or has incorrect measurements.
5. The cognitive engine can not connect to the radio.
6. The operating system can not connect to the cognitive engine.

The *induce* class of vulnerabilities is when the CR is stimulated to cause interference with a licensed transmitter. While the result is not an immediate DoS, it may cause permission policies to be tightened or eliminated potentially denying service over the long term. A CR may cause interference with a licensed transmitter under one or more of the following conditions:

1. The licensed spectrum appears unoccupied.
2. The policy is incorrect.
3. The location is incorrect.
4. The sensor provides incorrect measurements.
5. The commands to the TX/RX are incorrect.

These conditions parallel the DoS states except the 6th since, by design, no command from the operating system should induce the radio to transmit in an interfering channel. The attacks can be divided into broad areas that affect multiple vulnerabilities—such as a compromised cooperative CR—and attacks on specific vulnerabilities—such as the common control channel attacks etc.

In general there are six areas of security; confidentiality, privacy, integrity, authentication, authorization, and non-repudiation [24]. Confidentiality protects messages from being read by anyone but the intended recipient. Privacy protects the identity of the sender or receiver. Integrity prevents messages from being modified. Authentication validates the purported sender of a message to the receiver. Authorization controls access to services of authenticated users. Non-repudiation allows a receiver to prove that a message originated from its sender.

These areas as they relate to CR user traffic will not directly be considered. However, these techniques will be useful in preventing DoS attacks to the extent that they protect the signaling and communication between networked CR elements. Many mechanisms exist for these different techniques which we will assume in our discussion.

The standard approach to DoS is protection, detection, and reaction [8]. We should acknowledge here that since security is fraught with pitfalls that multiply with system complexity and require extensive system validation; the inherent complexity of CRs and the evolving system designs limit our discussion to general protection countermeasures rather than a complete solution. Such a solution is a part of ongoing and future work.

The subsequent sections describe in more detail these six avenues of attack, the relative effectiveness of each and the respective protection countermeasures.

A. *Spectrum occupancy failures*

A cognitive radio will not communicate on a channel that is being used by a licensed operator. A CR that detects such licensed use may be required to avoid the licensed channel for long periods of time. An attacker might mimic a licensed carrier. In this case there is a potentially large jamming gain for the attacker. Occupying each licensed channel for a brief time can prevent any channel from being used. For example, some CR will measure the channel it is using often. An attacker that detects a CR transmission can produce a signal with characteristics of the licensed transmitter until the CR radio detects the signal and ceases transmission. The attacker's signal need not be strong enough to physically jam the CR signal at the receiver. It only needs to be large enough to be detected by the CR transmitter.

This provides jamming gains in two dimensions. First, a short jamming period from the attacker can yield a long period of inactivity for the CR. CR can renegotiate new transmission bands if a licensed transmitter is detected. This typically takes time and some effort to negotiate a new channel between CR transmitter and receivers. Second, the CR's signal power and the attacker's signal power are independent of each other. The attacker's signal can be many orders of magnitude weaker than the CR's signal and yet still be detectable and thus prevent the CR transmitter from communicating.

Interestingly, if the attacker is near the transmitter, the power required to generate a detectable false signal can be below existing Part 15 limits so that it is possible the attacker's behavior does not directly violate any regulations.

Alternatively, the attacker may try to mask licensed users so that the CR will mistakenly communicate. In one approach the attacker may broadcast noise which raises the noise floor so that feature detectors tuned to the licensed service would fail. However, if the attacker generates too much noise, other more general power detectors will trigger. The level of noise power in order to mask the licensed transmitter may be low, and for some cognitive radios, the intervals when measurements are made are known and only a fraction of the total time. Together the average power for this attack may be low. However, the attack is fragile in the sense that strong licensed signals can not be masked in this way and it must mask the signal on every detection attempt to be successful. Further, if a CR is cooperating with other CR radios then every CR radio must be masked in this way.

These attacks where the attacker spoofs or masks a licensed transmitter are best dealt with from a cooperative architecture. With cooperating users, the attacker would need to appear as a licensed user to multiple CR that may be widely distributed which reduces the effectiveness of the attack. Alternatively if licensed user occupancy is well documented in a database or in information distributed via a broadcast beacon then a CR can use its location information and this database to have a reliable model of what white space is available. A non-cooperating user is more susceptible to this attack. It could also rely on a licensed user data base. In the case that a good data base is not available and only partial information is available, it has been shown that underlay-based schemes are more reliable than overlay schemes at avoiding interference to licensed users [16].

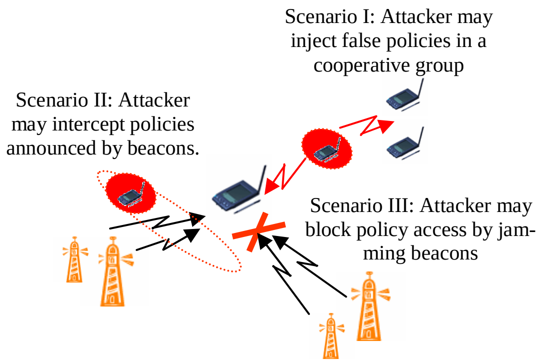
Scenario I: Attacker may inject false policies in a cooperative group
Scenario II: Attacker may intercept policies announced by beacons.
Scenario III: Attacker may block policy access by jamming beacons

Figure 3: Policy Failure Scenarios

B. *Policy failures*

A cognitive radio requires some policy that permits it to communicate. Policy failures include the lack of any policy or the use of false policies.

If the CR can be prevented from receiving any policy, then it will not communicate. This effect is more difficult to achieve. The policy database is not necessarily monolithic. The CR may draw on some general policies provided at time of manufacture (e.g. for unlicensed operation). A radio beacon may announce local policies. The CR may be able to make specific queries to a remote policy database. Or, the CR may be able to transfer policies from other CR. Furthermore, policies can be distributed in the form of certificates with a period of validity [10]. A CR may already have such certificates from an earlier access to the database.

At any given time, several policy certificates may be valid. These policies can be positive (permitting communication) or negative (preventing communications) and include conditions under which they apply. The cognitive engine must reason through these to find a sufficient policy for its intended communication or scale back its communication. An attacker can try to inject false policies into the CR policy database. Negative policies will prevent communication; positive policies may cause the radio to cause interference. Policies are introduced at the time of device manufacture, when the CR is updated, through general policy beacons, from other CR radios, and in response to queries to trusted policy databases. Each of these mechanisms, if used, is an opportunity to introduce false policies or modify valid policies. Figure 3 summarizes a few of the policy failure scenarios.

False policies can be prevented by having authentication and integrity certificates traceable to a trusted authority associated with each policy. These policies would have a lifetime associated with them that is preferably as long as possible. In this way policies can be freely exchanged among cooperative nodes and for non-cooperative nodes they would only require infrequent policy updates and renewals. [2]

By making it so that valid policies can be exchanged freely and with confidence and stored for long periods of time, it is unlikely that an attacker can prevent a CR from having at least some policies available.

C. *Location failures*

Almost all policies require some location reference. Even when using a pure sensing strategy, the CR must as a minimum know in which country or region it is operating in order to know which regulatory policy regime to apply. A greater number of policies can be applied as more specific location information is available. For instance, every TV channel is used somewhere in the United States. So, some other communication band must be used unless location information more specific than the country is available. Knowing one is in a specific region may create some opportunities. However, in some regions, such as New York City, the location must be known accurately, to within a few kilometers for the CR to be sure it will not interfere with any of the many TV transmitters in the New York metropolitan area [22]. From this example, it should be clear that any location information is useful. But, any degradation in location accuracy can limit or prevent communication.

Location information can come from standard geolocation techniques such as GPS or LORAN; user input of country, zip code, or street address; identifying known radio sources like FM radio or TV signals and finding their transmitter locations in a database; or from known location beacons such as from some cellular system base stations that broadcast their GPS coordinates. This diversity of sources enables a CR to always have some level of location awareness.

However, many of these sources are vulnerable. GPS signals are weak and easily jammed [9]. GPS often fails in indoor, dense urban or rough terrain environments. Manual entry is open to misconfiguration by intended users or malicious entry by users with access to the user interface. If the attacker is in close proximity, false TV signals can be generated that can be stronger than other TV signals. These attacks can cause the CR to have incorrect location estimates or increase the uncertainty in its estimates, both of which are effective at reducing the location specificity [20].

If the geolocator is networked, then the CR is gaining location information from outside sources such as a locator beacon or it infers its location from other CR radios. Attackers can generate false reports purported to be from these sources. Or, it can try to compromise these sources using one of the above techniques.

The key to having at least some location information available is for the CR to have multiple strategies for determining its location, especially if the CR is mobile. If it is cooperating it can share information with other users. In a centralized scheme with a subscriber base, subscriber nodes may be slaves to a trusted central authority and will only transmit under the permission and guidance of the central authority.

D. *Sensor Failures*

Section A already described failures that could be caused by false or masking inputs provided to sensors. A malfunctioning sensor could simply report false inputs with similar consequences. An attacker might also try to generate false reports purportedly from legitimate sensors.

If the cognitive engine can be prevented from receiving any sensor information then it will limit the communication options

[2] Of course, a part of the policy is its geographic application area and other requirements so that inappropriate policies would not be arbitrarily applied.

for the CR as many policies will require sensor measurements in order for them to be invoked. Sensor information exchanged via a common control channel provides a single and perhaps easy to jam channel.

In some CR designs, the sensor and radio share the same front end. Even when they are separate, the sensor sensitivity can be impaired by a nearby transmitter. As a result, sensing and transmission can not occur at the same time. The radio can only operate for some fraction of the time, *f*, with the remaining time being used for sensing. In this case, any jamming becomes leveraged by a factor of 1/*f*. For instance, if, because of sensing, the radio can only operate for *f* = 70% of the time. Then jamming 35% of the time will reduce the time for communication by 35%/*f* = 50%.

The key to avoiding leveraged jamming is to make the fraction of time devoted to transmission, *f*, as close to one as possible. Good sensing strategies are needed for this.

E. Transmitter/Receiver Failures

The receiver of a cognitive radio is often designed to work at a wider range of frequencies than typical radios. The antenna and receiver front end are therefore less selective. The receiver front end is potentially more susceptible to direct physical jamming that do not jam the signals directly, but instead seek to overload (desensitize) the front-end.

Different frequencies have different propagation characteristics. Jamming only lower frequencies may be sufficient to prevent communication. The CR may have available white spaces at higher frequencies but the propagation losses at these frequencies are too high to be useful.

Receiver errors may be perceived as evidence of licensed operation in the same band. In this case, jamming a receiver can cause the CR to abandon the band.

A key CR operation is for a transmitter and receiver to find each other to initiate communication. In a CR, the available frequencies depend on time and place so that some type of spectrum initiation protocol is needed. Once initiated, communication may need to change the frequency of operation due to the appearance of a licensed user or CR mobility, a so-called spectrum handoff [8]. These times are vulnerable because a failed initiation or handoff may require a long time for the radio to resume communication. An attacker can either induce a spectrum handoff via means described above, or recognize the CR signaling of a spectrum initiation/handoff and then start more aggressive jamming to cause a communication failure. By jamming only at these critical moments, the attacker has the potential to achieve a larger jamming gain.

In a networked CR an attacker that can gain control of the transceiver can prevent its use. Such an attack would be possible with any networked radio. However, with a CR the attacker could cause the radio to transmit and interfere with licensed users. It also opens the possibility of so-called Sybil attacks where the radio transmits using multiple identities, some of which behave while others misbehave [12].

As a countermeasure, the physical front end of a CR receiver needs to be designed for potentially large interfering signals. Use of multiple antennas or steerable antennas can enable the receiver to focus on the intended transmitter (and vice versa). Multiple antennas designed for different frequencies can also help mitigate the variable influence of frequency. The receiver also needs to be careful in how it interprets errors and use it as only one piece of evidence that there is a licensed user. As with the common control channel, the channel used for spectrum initiation/handoff has to be very robust and simple.

F. Operating System Disconnect

This attack can only be made if the cognitive radio is remote from the end user applications. Attacking this link would be the same whether the radio is cognitive or not and so is outside the scope of this paper. However, if the location information is to be provided via user input, then this disconnect does represent a new vulnerability, especially if this information can be selectively targeted.

In a distributed cooperative model the CRs may form an ad hoc or mesh network to distribute sensing and other information. Such networking is beyond simple physical or link layer access and so may require operating system support. A CR with disconnected or compromised operating system can inhibit or corrupt information dissemination and in general is subject to the well-known attacks on ad hoc networking [11].

A CR operating as an underlay will attempt to transmit so that its transmission does not cause interference to the licensed user. An attacker could add additional transmit power to the CR which might collectively cause interference.

G. Compromised cooperative CR

In a cooperative CR system, compromised nodes can be particularly insidious. They can produce false sensor information, false geolocation information, and invalid policies. They can also inhibit the forwarding and dissemination of valid information among CR nodes in a cooperative network. Authentication and integrity checks can mitigate the corruption of one user's data by others. In a centralized architecture, the central authority requires a public-key authentication and digital signature mechanism so that client CR can validate the source and integrity of the information. In the other direction, the central authority needs to also authenticate the source and integrity of the information. In some CR service models (e.g. if the secondary spectrum use is licensed), the client CR would be known subscribers and as in cellular, secret keys for each subscriber could be maintained by the central authority. Verifying distributed cooperative users would be more difficult.

A compromised user could still originate corrupted data. One possibility is to have "black-box" sensors (or geo-locators) that report measurements with their own authentication and integrity check. A public-key scheme could allow any user to recover the measurements from a known class of sensors and prevent any intermediaries, including a compromised CR associated with the sensor from corrupting the data. If time stamps are included with the data, then replay attacks would be avoided. Public keys could be distributed by certified authorities.

A compromised user can avoid forwarding sensor information. This falls in the realm of ad hoc network security issues which have been dealt with elsewhere [11]. In general the approach is for nearby nodes to identify a misbehaving node and then isolate it, for instance refusing to accept messages or otherwise interact with the isolated node. An inherent tradeoff is that a CR has more capable sensing to identify the compromised and misbehaving nodes. However, a compromised node has a more capable transmitter for masking its activities.

We note that when CR are non-cooperating, the value of a compromised CR is minimized to its local activity which can not be leveraged to more widespread disruption i.e. any attack with the compromised CR could have been performed with other radio types.

H. Common control channel attacks

A common control channel is a target for DoS attacks since successful jamming of this one channel may prevent or hinder all communication. For this reason, the channel should use a robust spread spectrum coding. The media access scheme should be robust and provide fair access. A complex media access protocol is 802.11. In 802.11 a number of unintended interactions between different elements and different layers have emerged that yield significant unfairness [14, 18]. Furthermore, such a complex media access protocol provides additional opportunities for attack [21]. Thus, fairness has to be thought through across multiple layers and the simplest access scheme focused on the control channel need is preferable.

VI. A Multi-Dimensional Analysis

While an attacker may mask a licensed user from a CR in a non-cooperative network architecture, it may not be able to institute the same attack in a cooperative setup. Similarly, the efficacy of an attack may vary according to the spectrum access technique and the spectrum awareness model applied by the victim CR(s). This section provides a qualitative analysis of how vulnerable are victim CRs to a potential DoS attack depending on these three different dimensions—the network architecture employed, the spectrum access technique used and the spectrum awareness model. Quantitative analysis to measure the effectiveness of the outlined attacks is part of ongoing and future work. No single metric is sufficient to assess the effectiveness of different attack scenarios. A number of metrics such as attack scenarios, attacker's resources to mount an attack and metrics to measure the effectiveness of the attack such as jamming gain [21], jamming efficiency [15], Packet Send Ratio and Packet Delivery Ratio [25] are currently under consideration. Nevertheless, the qualitative analysis here provides useful insights and design guidelines.

The CR network architecture determines how a CR mode of operation is vulnerable to attacks. A CR operating in non-cooperative network architecture has the advantage that more of its functionality is collocated and so can not be intercepted or jammed. However, it is more vulnerable to attacks that leverage the standalone operation of the device. Similar attacks are more difficult to be successful when launched in a cooperative CR network, as the member CRs can validate the network measurements against each other in a distributed setup, or a central authority can validate measurements received from the CR group in a centralized setup. In other words, since a distributed cooperative network setup allows CRs to collate information about their radio environment, it provides an inherent security against device-centric DoS attacks. For instance, an attacker that emulates a licensed user can easily deny communication to a victim CR which is using detection as its spectrum sensing approach in a non-cooperative network as against one in cooperative setup. Similar analogy can be applied to injecting false policy, location or sensor information attacks. These attacks are easier to implant in CRs that rely on detecting these inputs without validating them with their neighbors as in a non-cooperative setup. On the contrary, attacks to intercept sensitive information, such as operating carrier frequency of different member CRs in a cooperative group, are easier to launch in a cooperative network setup where a malicious member CR can listen to exchange of such sensitive information over common control channels.

The access methods determine how the transmitted signal is vulnerable to the different attacks. As noted earlier a CR can generally operate in many frequency bands and so has inherent frequency diversity protection against direct DoS attacks. However, the overlay and underlay are not identical in how they realize this diversity. Generally, the underlay scheme, which spreads its spectrum over a large swath of bandwidth, is not as vulnerable to attacks which attempt to induce the CR to communicate in a licensed band. The power transmitted in any one band is low and so errors in identifying primary users have less effect. Since the underlay scheme has less interference with primary users it can afford less frequent access to sensor information and is less likely to need sudden spectrum handoffs as the spectrum environment changes. The wideband underlay scheme has an inherently lower vulnerability to direct jamming of data and control information. Against these many security advantages, the underlay radio is generally more complex to implement and potentially requires multiple complex filters to notch out critical bands (e.g. aviation radars). Alternatively, the underlay scheme can operate over a smaller bandwidth; however this limits the communication range.

The spectrum awareness method determines how the information used by the cognitive radio to make its spectrum selection is vulnerable to attack. The beacon method requires only one way interaction. In general, the beacon source already knows the available channels so that it requires no direct sensing or location information and, in effect, distributes the policy directly. Further, the beacon can make planned changes to spectrum usage that require less peer-to-peer coordinated spectrum handoffs. However, the beacon is a single point of failure and an attacker may be able to jam beacon information, inject false beacon information, or use the beacon information to predict CR activity. A further deficiency is that the beacon method may be too coarse a resolution if there are few beacons (e.g. one per metropolitan area).

The method that geolocates itself and then accesses a database can have finer grained policies. It is also independent of sensing like the beacon method, but depends on reliable location information. The radio may be able to store multiple valid policy certificates that depend on location. Individual policy queries which could be encrypted would need to be intercepted individually to predict CR activity. As with beacons, spectrum handoffs can be anticipated and planned ahead of time. This method is vulnerable to attacks on its location data and its access to the remote policy database. Both of these are likely to require cooperative sharing on location data or access to the database and so are more vulnerable to attackers that do not forward information.

Sensor based detection has the main advantage that it can operate independent of supporting infrastructure and does not necessarily depend on cooperation with other nodes. However, it does depend on sensor data, is required to spend time making sensor measurements, and may need to make sudden handoffs when a licensed user appears.

Based on this discussion, a summary of the relative vulnerabilities along each dimension are shown in Table 1. For each dimension and each vulnerability a “+” indicates the design choice is less vulnerable (more secure), a “−” indicates the design choice is more vulnerable (less secure), and a “.” indicates that that it is not significantly better or worse. Though somewhat of an oversimplified view, it provides a means to readily compare the choices along each dimension. This view should also be considered against the threat posed by the individual attacks. Finally, only DoS attacks are reviewed in this table and not other design factors such as cost or flexibility.

From a security perspective, the underlay approach appears to be the clear choice. The cooperative network architectures do better than non-cooperative with the centralized approach the better of the two cooperative architectures. For the spectrum awareness methods, the beacon and the geolocation methods fare similarly well. However both require an investment in infrastructure that is not necessary for the detection method.

VII. CONCLUSION

A naïve cognitive radio design will be vulnerable to multiple modes of failure from intentional and unintentional attacks. Any radio is subject to direct jamming. The cognitive-radio-specific attacks differ in that they have high jamming gain: they can induce large performance degradation for relatively little effort in terms of average transmission power. However, with modest effort on the part of the cognitive radio design, these jamming gains can be significantly reduced. Moreover the multi-dimensional analysis provided in the paper highlights certain inherent risks exposed and certain others minimized by the CR network architecture employed, the spectrum access technique used and the spectrum awareness method used by victim CRs. Cognitive radio designers are encouraged to consider these assessments in the light of potential vulnerabilities and remedies as they continue to develop cognitive radios.

REFERENCES

[1] A. Householder, A. Manion, L Pesante and G. M. Weaver, “Managing the threat of denial-of-service attacks,” CERT Coordination Center, v10.0, October 2001.

[2] DARPA XG Working Group, “The XG vision,” Request for Comments, version 1.0, Prepared by BBN Technologies, Cambridge, Mass., USA. July 2003.

[3] DARPA XG Working Group, “The XG architectural framework,” Request for Comments, version 1.0, Prepared by BBN Technologies, Cambridge, Mass., USA. July 2003.

[4] End-to-End Reconfigurability (E2R) Phase II website (e2r2.motlabs.com).

[5] FCC ET Docket No. 02-135, “Spectrum policy task force report,” Nov. 2002. (http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-228542A1.pdf).

[6] FCC ET Docket No. 03-108, “Facilitating opportunities for flexible, efficient, and reliable spectrum use employing cognitive radio technologies,” FCC Report and Order adopted on March 10, 2005, (http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6517509341).

[7] FCC, ET Docket No. 04-186, “Unlicensed operation in the TV broadcast bands,” ET Docket No. 02-380, “Additional spectrum for unlicensed devices below 900 mhz and in the 3 ghz band,” FCC Report and Order And Further Notice of Proposed Rulemaking, adopted on October12, 2006.

[8] I. F. Akyildiz, W. Y. Lee, M. C. Vuran, S. Mohanty, “NeXt generation dynamic spectrum access cognitive radio wireless networks: A survey,” Computer Networks, 50, 2006 pp. 2127-2159.

[9] J. A. Volpe, “Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System,” Final Report for the National Transportation Systems Center, U.S. Dept. of Trans., Aug. 29, 2001 (www.navcen.uscg.gov/gps/geninfo/pressrelease.htm).

[10] J. M. Chapin, W. H. Lehr, “Time-limited leases for innovative radios,” in Proc. IEEE DySPAN, April 17-20, 2007.

[11] J. Hubaux, L. Buttyán and S. Čapkun, “The quest for security in mobile ad hoc networks,” ACM press, New York, NY, 2001.

[12] J. Newsome, E. Shi, D. Song, A. Perrig, “The Sybil attack in sensor networks: analysis & defenses,” in Proc. of Third International Symposium on Information Processing in Sensor Networks, IPSN, 26-27 Apr. 2004, pp. 259-268.

[13] L. Ma, X. Han, C.-C. Shen, “Dynamic open spectrum sharing MAC protocol for wireless ad hoc network,” in Proc. IEEE DySpan, Nov. 8-11, 2005, pp. 203-213.

[14] M. Heusse, F. Rousseau, G. Berger-Sabbatel, A. Duda, “Performance anomaly of 802.11b,” in Proc. of INFOCOM 2003, v. 2, 30 Mar. - 3 Apr. 2003, pp. 836-843.

[15] M. Ståhlberg, “Radio jamming against two popular mobile networks,” 2000, Helsinki University of Technology, Tik-110.501 Seminar on Network Security.

[16] R. Menon, R. M. Buehrer, J. H. Reed, “Outage probability based comparison of underlay and overlay spectrum sharing techniques,” in Proc. IEEE DySpan, Nov. 8-11, 2005, pp. 101-109.

[17] R. Shirey, “RFC 2828: Internet Security Glossary,” IETF, May 2000.

[18] S. Pilosof, R. Ramjee, D. Raz, Y. Shavitt, P. Sinha, “Understanding TCP fairness over a wireless LAN,” in Proc. of INFOCOM 2003 v. 2, 30 Mar. - 3 Apr. 2003, pp. 863-872.

[19] S. Sankaranarayanan, P. Papadimitratos, A. Mishra, S. Hershey, “A bandwidth sharing approach to improve licensed spectrum utilization,” in Proc. IEEE DySPAN, Nov. 8-11, 2005, pp. 279-288.

[20] T. X. Brown, “An analysis of licensed channel avoidance strategies for unlicensed devices,” in Proc. IEEE DySPAN, Nov. 8-11, 2005.

[21] T. X. Brown, J. E. James, A. Sethi, “Jamming and sensing of encrypted wireless ad hoc networks,” in Proc. Seventh ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), Florence, 22-25 May 2006.

[22] T. X. Brown, D. Sicker, “Can cognitive radio support broadband wireless access?,” in Proc. IEEE DySPAN, April 17-20, 2007.

[23] T. X. Brown, A. Sethi, “Potential cognitive radio denial of service attacks and remedies,” in Proc. International Symposium on Advanced Radio Technologies 2007 (ISART 2007), Boulder, 26-28 Feb 2007.

[24] W. Stallings, “Network security essentials: applications and standards,” 3rd Ed., Prentice Hall, 2006, pp. 432.

[25] X. Wenyuan, T. Wade, Z. Yanyong, W. Timothy, “The feasibility of launching and detecting jamming attacks in wireless networks,” in Proc. ACM International Symposium on Mobile Ad Hoc Networking and Computing, Illinois, 2005, pp. 46- 57.

Table 1: Multi-dimensional Analysis of CR-specific DoS Attacks (+ is more secure).

Attack	CR Network Architecture			Access Method		Spectrum Awareness Method		
	Non-cooperative	*Centralized Cooperative*	*Distributed Cooperative*	*Overlay*	*Underlay*	*Beacons*	*Access DB / Geolocate*	*Detection*
Attacker emulates licensed user. Exploits long time before licensed channel can be reused	−	+	+	.	.	+	+	−
Attacker masks licensed user	−	+	+	−	+	+	+	−
Attacker blocks access to policies	.	.	+	.	.	−	.	+
Attacker injects false policies	−	+	−	−	+	−	.	.
Attacker intercepts policy information to predict CR activity.	+	−	−	−	+	−	+	.
Attacker blocks access to location information	−	+	.	.	.	+	−	.
Attacker blocks access to sensor information	+	−	.	.	.	+	+	−
Attacker leverages jamming against fraction of time transmitting versus sensing	−	+	.	−	+	+	+	−
Attacker induces receiver errors as if from a primary device	−	+	+	−	+	+	+	−
Attacker jams at spectrum handoff or initiation	.	.	.	−	+	+	+	−
Attacker misbehaves forwarding information between networked CR	+	+	−	.	.	.	−	+