Mobile Netw Appl (2008) 13:516–532
DOI 10.1007/s11036-008-0085-x

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

Timothy X Brown · Amita Sethi

Published online: 13 July 2008

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

1 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 [1]. 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 [13]. A number of research efforts such as the Defense Advanced Research Projects Agency Next Generation project [9, 10], the IEEE 802.22 Working Group [19] in the United States, and the End-to-End Reconfigurability program in Europe [11] 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 no other advance "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" [14].

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 two earlier papers [5, 6] but refines the analysis to assess

T. X Brown (✉)
Interdisciplinary Telecommunications,
Electrical and Computer Engineering, University of Colorado,
Boulder, CO 80309, USA
e-mail: timxb@colorado.edu

A. Sethi
Interdisciplinary Telecommunications,
University of Colorado,
Boulder, CO 80309, USA
e-mail: sethi@colorado.edu

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.

2 Traditional versus cognitive radios

While a traditional radio allows minimal user interaction and has unalterable receiver transmitter operations, the CR houses advanced functionalities of [1]:

- *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 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.

3 Operational cognitive radio aspects

3.1 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

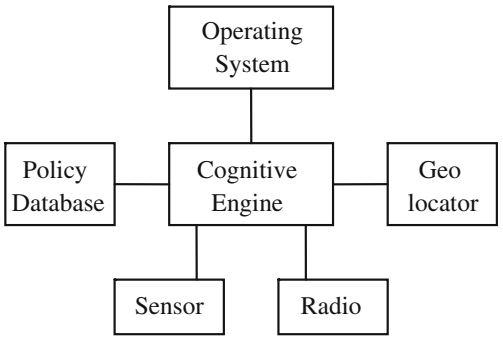

Figure 1 Cognitive radio components

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.

3.2 Architectures

CRs can be broadly classified into one of three network architectures as shown in Fig. 2. 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, geo-location, or communication extensions. Many cooperative schemes (centralized or distributed) envision a common control channel that is a well known link to share information [20, 25]. The operational advantages or disadvantages for any architecture in this range are not considered. Rather, we consider the security vulnerabilities when any of the cognitive radio components are or are not collocated with each other.

3.3 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

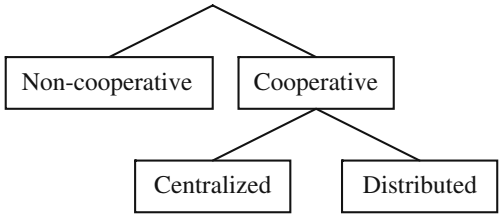

Figure 2 Cognitive radio network architectures

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.

3.4 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 [15]. A CR may use the

- *geo-location/database* approach to geo-locate itself and then 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 coverage area of the beacon.
- *detection/sensing* approach to directly measure radio activity, 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.

4 Denial-of-service vulnerabilities

A denial-of-service attack is an act of preventing authorized access to a system resource or the delaying of system operations and functions [26]. 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.

4.1 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* [4]. 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 for how it networks with other CRs; the spectrum access technique it uses; and the spectrum awareness model—open up jamming gains to the attacker.

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

4.2 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 [4]. 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.

4.3 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 signals 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.

How close must the attacker be to the victim? The power needed to institute the attack decreases as the attacker comes closer. As the distance decreases 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 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.

5 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.
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 [27]. 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 [17]. 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.

5.1 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 [22].

5.2 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 [8]. 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.

5.3 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

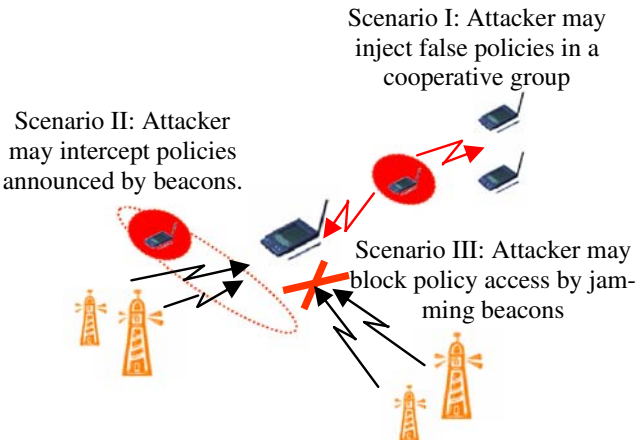

Figure 3 Policy failure scenarios

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 [7]. 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, postal 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 [29]. 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 [3].

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

[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.

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.

5.4 Sensor failures

Section 5.1 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 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.

5.5 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. However, the attacker must be close to the cognitive radio or use high jamming power to be effective.

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 [1]. 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 [23].

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.

5.6 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 [18].

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.

5.7 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 geolocators) 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 [18]. 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.

5.8 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 [16, 24]. Furthermore, such a complex media access protocol provides additional opportunities for attack [4]. Thus, fairness has to be thought through across multiple layers and the simplest access scheme focused on the control channel need is preferable.

6 A multi-dimensional analysis

While an attacker may mask a licensed user from a CR in 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 presents a qualitative analysis of the potential CR denial-of-service attacks in terms of the likelihood of the attack's occurrence, its impact on victim CRs and the overall risk posed by the attack. A set of 18 network architecture, spectrum access technique, and spectrum awareness combinations are defined as well as 12 exemplary attacks. These are analyzed according to the methodology described in [12], also used in [2]. This attack evaluation scheme provides a qualitative assessment of the risk posed by an attack. While these risk evaluations are subjective and specific to the system design combinations and attacks defined, they provide some useful insights into the types of risks and the relative risk of different system design choices.

Quantitative analysis to measure the effectiveness of the outlined attacks is part of ongoing and future work. A number of metrics such as attacker's resources to mount an attack and metrics to measure the effectiveness of the attack such as jamming gain [4], jamming efficiency [28], Packet Send Ratio and Packet Delivery Ratio [30] are currently under consideration. However, the qualitative analysis allows us to compare a larger set of design choices and a larger number of attacks.

6.1 Analysis method

The analysis is based on the following three criteria, the last of which is derived from the first two:

- *Likelihood* evaluates the possibility of an attack's occurrence. It is measured by evaluating the technical difficulty faced by the attacker while conducting the attack. The technical difficulty is strong if it requires significant resources or there are a number of unknowns; solvable if it requires modest effort and the information required to conduct the attack is available

or weak if there is no technical difficulty involved in conducting the attack. Accordingly the attack's likelihood varies from low, possible, to likely (see Table 1). For instance any attack that requires cryptanalysis of encrypted or digitally signed data is considered to have strong technical difficulty and thus low likelihood.

- *Impact* evaluates the attack's affect on victim CRs communication. The attack's affect on performance may range from mere annoyance; to noticeable yet still operational; to completely non-operational CR communication. Accordingly the impact varies from low, medium to high. Table 2 succinctly defines the levels of impact for both denial and induce attacks.
- *Risk* posed by an attack is calculated from the product of rank values assigned to the likelihood and impact. A one or two risk value indicates minimal risk to victim CR and requires no countermeasures. A three of four risk value indicates major risk to the victim CR and indicates that the threat cannot be ignored. Finally, a six or nine risk value indicates a critical threat that needs to be handled at high priority,

The attack evaluation methodology is summarized in Tables 1, 2 and 3. The likelihood and impact values are assigned based on subjective assignment by the authors.

6.2 System design space

The efficacy of an attack may vary according to the network architecture, spectrum access technique and the spectrum awareness model applied by the victim CR(s). The CR network architecture determines how a CR is vulnerable to attacks. A CR operating in a *non-cooperative* network architecture has the advantage that more of its functionality is collocated, i.e., does not depend on networked information exchange with peers, and so the elements of its CR operation 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 cooperative group members can collate network measurements against each other as peers in a *distributed cooperative* architecture, or a central authority can validate measurements received from client CR nodes and override with its own in a *centralized cooperative* architecture. In other words, a cooperative network setup allows CRs to collate information about their radio environment, providing an inherent security against device-centric DoS attacks.

The spectrum access technique is either *underlay* or *overlay* and refers to the technique used by the CR to communicate with peer CR. 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. 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). The underlay scheme can operate over a smaller bandwidth without increasing its power in any individual primary user channel; 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. We explain each of the three in detail. The *beacon* method requires only one way interaction. In general, the beacon source already

Table 2 Summary of the impact component of analysis

Rank	Cases	Rationale: denial attacks	Rationale: induce attacks
1	Low	Perceptible but insignificant degradation in CR communication	Perceptible but infrequent interference to active primary users
2	Medium	Significant degradation but still operational CR communication	Perceptible frequent interference to active primary users
3	High	Non-operational CR communication	Continuous interference to active primary users

Table 1 Summary of the likelihood component of analysis

Rank	Cases	Rationale: technical difficulty
1	Low	Strong
2	Possible	Solvable
3	Likely	None

Table 3 Summary of the risk component of analysis

Rank	Cases	Rationale: technical difficulty
1,2	Minor	No countermeasures required
3,4	Major	Threat cannot be ignored
6,9	Critical	Mandates high-priority handling

knows its location and 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 also 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.)

In the *geolocation and database* method the CR geolocates itself and then compares its location to a database of policies as a function of location. It is 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. Cooperative techniques provide alternate sources of location and policy information.

Sensor based detection has the main advantage that it can operate independent of supporting infrastructure and does not necessarily require outside communication (with beacon, geolocation, or policy sources). 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.

In this analysis we do not specifically compare different types of operation such as low-power WLAN or high-power WMAN applications. Nor do we compare mobile vs. fixed applications. Such comparisons are part of future analysis.

6.3 System design choices

We consider 18 combinations of three network architectures, two spectrum access techniques and three spectrum awareness models. In our risk analysis we found that the effect of spectrum access was isolated to its consequence on attacks that induced communication in occupied bands or induced spectrum handoffs. In both cases the impact on the underlay scheme was less than the overlay scheme. However, the nine combinations of network architecture and spectrum access technique could have a variety of interrelated interpretations so we define them precisely for each combination based on the general definitions provided earlier.

Beacon, non-cooperative In this design each CR gathers information about available spectrum from beacons. The channel information is derived from a central database. The database update mechanism and how it is connected to the beacon transmitter are outside the scope of this attack analysis. The beacon signal is one way from the beacon transmitter to the CR and is unencrypted. The beacon information uses a certification scheme so that spoofing messages would be difficult. Timestamps and/or sequence numbers protected by digital signatures prevent replay attacks. Location information is also included in the message to deter forwarding of beacon messages to different areas. The location in the message compared to the location of different beacons in a region can be checked for consistency. The policy information consists of time-limited certificates valid less than an hour. The isolated CR does not have any sensing or geolocation capability. The beacon information is relevant to the CR by virtue of its being received by the CR.

Beacon, distributed cooperative This design uses the same beacon as in the non-cooperative case. However, CR are allowed to forward and share beacon information. These CR communicate using ad hoc networking, however the routing includes hop-count information so that a CR can track the relevance of forwarded information it receives. The information CR exchange is encrypted.

Beacon, centralized cooperative In this design, client CR are commanded by a centralized radio source. These commands use a common control channel. The common control channel architecture is two way but independent of the underlay/overlay scheme used for communication between CR. The CR authenticates and registers with the central controller. The CR can then get spectrum policy information that is encrypted and unique to that CR. Symmetric key rather than public key methods can be used to encrypt, certify, and sign the information. The policy information consists of time-limited certificates valid less than an hour. The CR user traffic may be between the CR and the central controller (like in a cellular scheme) or it may be between peer CR (the central controller is only a source of policy information on available spectrum).

Geolocation and database, non-cooperative In this design individual CR locate themselves using GPS. Their location is compared to an internal database of policies. These policies indicate what channels can be used relative to a location. The policy information consists of time-limited certificates valid for periods of at least a day. The CR would need to occasionally query a central database to maintain a set of valid policies for its current location. These queries would be authenticated and encrypted. The policies could be finer spatial resolution corresponding to the accuracy of

the location method and how often the location is updated. For frequent updates (at least once a minute) with GPS the policy could be unique to the current square kilometer. Less frequent updates would require policies valid over larger areas. The CR does not have a beacon receiver or sensing capability.

Geolocation and database, distributed cooperative This design uses the same location and policy database mechanisms as the non-cooperative case. However, CR are allowed to forward and share location and policy information. The location information has an increasing error added with each additional transmission hop. The policy information is freely shared along with the original certificates so that it can be authenticated by the receiver. The information CR exchange is encrypted.

Geolocation and database, centralized cooperative This design has the same mechanisms as the distributed cooperative design, except that CR can share location and policy information only with a central controller. The central controller has reliable access to a policy database and provides a geographic reference point.

Detection and sensing, non-cooperative In this design individual CR sense the radio environment to maintain a database of spectrum that is apparently free or is occupied by primary users. The potential spectrum and the policies for sensing and access are maintained in an internal database. The policy information consists of long term time-limited certificates valid for periods of at least a month. The CR would need infrequent queries to a central database to maintain a valid set of policies. These queries would be authenticated and encrypted. The CR is aware of the country in which it is operating via user input. The CR does not have a beacon receiver.

Detection and sensing, distributed cooperative This design uses the same sensing and policy information as in the non-cooperative case. However, this information can be shared among CR. These CR communicate using ad hoc networking, however the routing includes hop-count information so that a CR can track the relevance of forwarded information. The policy information is freely shared along with the original certificates so that it can be authenticated by the receiver. The information CR exchange is encrypted. Having multiple sensors in an area increases the detection reliability since primary users are more likely to be detected if there are more sensors.

Detection and sensing, centralized cooperative This design has the same mechanisms as the distributed cooperative design, except that CR can share sensor and policy information only with a central controller. The central controller has reliable access to a policy database and can aggregate sensor information over a larger area. The central controller has a directional antenna or other method for providing some localization of the different CR and their sensor data. For instance a sectorized antenna would enable the central controller to limit the scope of sensor measurements to each sector.

From the above discussion, it can be concluded that the three CR dimensions are intertwined and thus, an attacker can exploit vulnerabilities exposed in each combination possible with the three different dimensions. A discussion of specific attacks is described next.

6.4 Attack description

The vulnerabilities described in Section 5 suggest a number of attacks that are possible. We analyze 12 different attacks using the risk assessment methodology. Table 4 summarizes the attacks and our risk assessment for the 12 attacks versus the 18 CR designs. For each combination, a square is placed whose height (1, 2, or 3) indicates the likelihood of the attack (i.e. how easy is it to execute) and width (1, 2, or 3) indicates its impact on the user if the attack is successful. No square indicates the attack is not relevant or has no impact on that design. We address each of the attacks in turn.

(1) Attacker emulates licensed user

An attacker that emulates a primary user can temporarily cause victim CRs, which use sensing to detect the presence of licensed users, to abandon the licensed channel irrespective of whether the CR is operating in cooperative or non-cooperative mode. The attack does not pose technical difficulties to the attacker till he is focused in a specific licensed band. It is, however, technically difficult to launch the attack in every possible channel so as to completely deny CR communication. The difficulty is even more for cooperative architectures as the cooperative group members can collate the sensor measurements with peers or with a central entity and conclude that the primary user signal is fake. However, if successful the attack can deny CR to communicate, thus has a high impact ranking. This attack poses the same risk irrespective of whether the victim CR transmits through overlay or underlay spectrum access.

(2) Attacker masks licensed user

An attacker can attempt to mask the presence of the primary user through low-level jamming that induces CR which use sensing to interfere with a primary user. However, this attack is only effective when the primary user signal is weak and the attacker carefully regulates its power so as to not trigger general power sensing. If

Table 4 Cognitive radio DoS attack risk assessment of 12 attacks vs. 18 design combinations

Attack	Overlay									Underlay								
	Beacon			Geolocation Database			Detection Sensing			Beacon			Geolocation Database			Detection Sensing		
The attacker …	NC	CC	DC	NC	CC	DC	NC	CC	DC	NC	CC	DC	NC	CC	DC	NC	CC	DC
emulates licensed user							■	■	■							■	■	■
masks licensed user							■	■	■							■		
blocks access to policies	■									■								
injects false negative policies	■	■	■	■	■	■				■	■	■	■	■	■			
injects false positive policies	■	■	■	■	■	■				■	■	■	■	■				
intercepts policy information to predict CR activity	■		■	■		■				■		■	■		■			
blocks access to location information				■	■	■							■	■	■			
blocks access to networked sensor information								■	■								■	■
leverages jamming against fraction of time transmitting vs. sensing							■	■	■							■	■	■
induces receiver errors as if from primary device							■	■	■							■	■	■
jams at spectrum handoff or initiation	■		■				■	■	■	■		■				■	■	■
misbehaves forwarding information between networked CR			■			■			■			■			■			■

successful it will add an additional channel to the available spectrum for the CR that may or may not be used. The technical difficulty of this task is high and the impact is medium for non-cooperative detection-based designs. The impact for cooperative schemes is lower since masking is difficult to simultaneously achieve for multiple CR and one of the cooperative CR is likely to detect the primary user eventually. The underlay spectrum access scheme has lower impact since its power is spread out and it is unlikely that a small amount of power radiated into a primary user's band will cause harmful interference.

(3) Attacker blocks access to policies

The CR requires at least one policy that gives it permission to communicate in some band. An attacker can attempt to block access to the messages carrying these policies. In the beacon spectrum awareness method the policy messages are sent in the clear in the non-cooperative and distributed cooperative architectures and an attacker can attempt to jam these at a specific CR. In the distributed cooperative, the attacker would need to jam the beacons and policy exchange messages from nearby CR as well. The latter can be encrypted among the members of the CR group. In the centralized control architecture, the policy messages are encrypted and more difficult to intercept. Since beacon policies are shorter lived the CR requires a consistent source of policy messages to continue communicating. Conversely the attacker would need to consistently jam the beacon messages in order to prevent communication. The geolocation and database policies are much longer lived. The attacker would need to diligent monitor the CR for days at a time in order to prevent the CR from obtaining any policy. With detection and sensing the policies are so long lived that blocking of policies is not relevant. If all policies were effectively exhausted, the CR would stop communicating or be limited to default minimal policies. The attack does not depend on whether the access method is overlay or underlay.

(4) Attacker injects false negative policies

As noted earlier, false negative policies will prevent CR communication. For the sake of analysis, we assume that all policy messages carry a digital signature and a suitable certification chain so that an attacker could not readily create false policies. Further, the policies have sequence numbers and time stamps so that replay attacks

would be ineffective. The timestamps would not need to be more accurate than to the nearest few seconds (i.e. millisecond clock synchronization accuracy is not needed). Finally, location context would prevent policies from one location to be replayed in another location. In a beacon scheme, the CR could compare the location context of other beacons it may have heard or be hearing to filter out beacons with incongruent location context. The geolocation and database approach has an independent source of location information. Further, it makes its own encrypted queries to a policy database. The detection and sensing makes rare policy update enquiries and further these enquiries are encrypted. As a result this attack is considered ineffective for detection and sensing. In the centralized cooperative architecture the messages are sent encrypted for individual CR and thus more difficult to spoof. In the distributed cooperative architecture, neighboring false beacons could be heard by all CR and thus comparing beacon information may not reduce the impact of the attack. However, with the geolocation and database method in a distributed cooperative setup, an attacker is unlikely to send false database response messages to every CR and so the impact is less. This attack does not depend on whether the access method is overlay or underlay.

(5) Attacker injects false positive policies

An attacker can also inject false positive policies to induce radio to communicate and thereby cause interference. This attack has all the issues as with injecting false negative policies with two exceptions. First, we assume that CRs are conservative so that if a radio holds a policy that prevents the use of a band and a policy that permits the use of a band, the preventing policy will prevail. In the case of geolocation and database, where the policies are long lived, a single false positive policy is unlikely to have an effect over existing long policies. The attacker would need to insert false policies over a long period in order to be effective. In short the attack has less impact in the geolocation and database case. A second exception is that false positives have inherently less impact with underlay schemes since only a small amount of additional power would be added to the permitted bands by an individually compromised radio.

(6) Attacker intercepts policy information to predict CR activity

An attacker that knows which bands a CR is permitted to operate can create more focused attack strategies. We assume that the main impact is that the attacker is more efficient in their attacks (i.e. level 1). Schemes that send encrypted policy information have the lowest likelihood. Unencrypted beacon schemes have the highest likelihood. Again the detection and sensing schemes have such infrequent access to policy information that this attack is considered irrelevant. These attacks do not depend on whether the access method is overlay or underlay.

(7) Attacker blocks access to location information

GPS signals are susceptible to jamming by an attacker. Without location information the geolocation method will fail. In the centralized cooperative scheme the centralized authority can always provide an approximate location that the CR can use. In the distributed cooperative scheme a nearby CR not subject to GPS jamming can provide an approximate geolocation. This attack does not apply to beacon and the detection and sensing spectrum awareness methods. This attack does not depend on whether spectrum access is overlay or underlay.

(8) Attacker blocks access to networked sensor information

A CR without access to sensor data can not detect primary users and if persistent will not satisfy the requirements of listen before talk policies. Thus this can induce or prevent communication. This attack only applies to centralized cooperative or distributed cooperative architectures that use detection and sensing. The centralized authority depends on spectrum measurements from client CR in order to make its spectrum assignment decisions. The attacker would need to block access to all of these measurements and further they are sent over encrypted links. In order to block on an encrypted link the attacker would need to brute force block all communication. The distributed cooperative uses the measurements of neighbors to improve the detection accuracy, however the CR can operate without any of these measurements. These messages are also exchanged encrypted. This attack would have less impact on underlay schemes in the case that false positive spectrum assignments are made.

(9) Attacker leverages jamming against fraction of time transmitting versus sensing

By carefully jamming only when the CR is sending an attacker can achieve a small jamming gain in the detection and sensing spectrum awareness method. In centralized cooperative architectures, the sensing load is distributed via encrypted commands from the centralized authority. In other architectures it may be somewhat difficult to predict when the CR is detecting vs. receiving or transmitting. If successful the impact is low. This attack does not depend on whether spectrum access is overlay or underlay.

(10) Attacker induces receiver errors as if from a primary device

In the detection and sensing spectrum awareness method, one sensing input is the state of the channel as perceived by the radio. Errors on this channel may be interpreted as the presence of primary radios. As a result the CR may perform a spectrum handoff or list the channel as occupied. This attack is of modest difficulty and has low impact for both overlay and underlay network access. Cooperative schemes can compare sensor information and would require a coordinated attack on multiple CR in order to be effective. This attack does not depend on whether spectrum access is overlay or underlay.

(11) Attacker jams at spectrum handoff or initiation

The attacker can jam at spectrum handoff which may cause the handoff to fail and the CR to engage in a longer rendezvous process. To be effective, the attacker should anticipate the handoff and jam information exchanged during the handoff process. The attacker can also jam immediately after the rendezvous so that the CR deems the rendezvous a failure and immediately initiates a new rendezvous process. Spectrum handoffs can occur when the spectrum policy changes, a time-limited policy expires, the CR enters a new location, or the sensing detects a change in the occupied channels. The spectrum handoff only occurs if the CR is using a channel that is not allowed under the new situation. With no prior information, an attacker can monitor a CR and when it vacates a channel attempt to detect it communicating on a new channel and then jam at that point. This is technically difficult, is unlikely to be achieved with certainty, and as a result, the likelihood is low and the net impact is medium. The likelihood of the attack increases if it has access to the policy changes such as with non-cooperative and distributed cooperative beacons. In the case of detection and sensing, the attacker can emulate primary users and thus control the timing of some spectrum handoffs. The emulates-licensed-user attack described earlier has lower likelihood of success because to be successful it must emulate a licensed user in a significant fraction of the channels. Here the attacker only needs to spoof the CR on one channel to initiate the spectrum handoff which is much easier for the detect and sense with non-cooperative architecture. For the cooperative schemes because of shared detection the likelihood is lower. Underlay schemes are inherently more robust to this attack since changing available spectrum results in adjustments to the distribution of their signal and not hard handoffs.

(12) Attacker misbehaves forwarding information between networked CR

These attacks require the attacker to compromise a CR node which we assume is likely. This attack affects the distributed cooperative network architecture. The information exchanged includes policy information, location information, and sensing information. Since policy information has digital signatures from their source, this information is unlikely to be considered valid, but the nodes will require extra work to filter out these messages. Conversely, the CR can not forward policy information reducing the information available to nodes. Occasionally nodes may be left with no valid policy as a result. False location information can be generated by a compromised node and forwarded to cooperative CR which will dilute their location accuracy. Finally, false sensor reports can be generated. A compromised node can falsely identify primary users in many bands and significantly limit communication. These attacks do not depend on whether the access method is overlay or underlay.

This attack analysis is summarized in Table 4. Every attack is assigned a risk box that is calibrated to show the likelihood rank on the vertical axis and impact rank on the horizontal axis. The total number of shaded squares contained in a box shows the overall risk posed by the attack in the specific CR configuration. This analysis is discussed in the next section.

6.5 Attack analysis and discussion

One goal of this paper is to provide cognitive radio design recommendations. Table 5 summarizes the total risk along each of the three dimensions and is derived from the data in Table 4. For instance, when an attacker emulates a licensed user; the total risk for the non-cooperative network architecture is 12 (sum of columns 1, 4, 7, 10, 13 and 16 in Table 4) compared to 6 for each of the cooperative architectures; the total risk for the overlay access method is 12 (sum of columns 1–9 in Table 4) compared to 12 for underlay; and the total risk for the detection and sensing access method is 24 (sum of columns 7–9 and 16–18 in Table 4) compared to 0 for the other methods. From such data we can conclude that for this attack, along the network architecture dimension non-cooperative is more vulnerable; along the spectrum access dimension overlay and underlay are equally vulnerable; and along the spectrum awareness dimension, detect and sense is the most vulnerable. Such a view can be applied to each of the attacks.

Table 5 Multi-dimensional analysis of CR-specific DoS attacks

Attack	CR network architecture			Access method		Spectrum awareness method			Total
The attacker ...	Non-cooperative	Centralized cooperative	Distributed cooperative	Overlay	Underlay	Beacons	Geolocate database	Detection sensing	
Emulates licensed user	12	6	6	12	12	0	0	24	24
Masks licensed user	3	1	1	4	1	0	0	5	5
Blocks access to policies	18	12	8	19	19	22	16	0	38
Injects false negative policies	12	12	10	17	17	18	16	0	34
Inserts false positive policies	8	8	6	14	8	15	17	0	22
Intercepts policy information to predict CR activity	8	4	8	10	10	14	6	0	20
Blocks access to location information	18	12	12	21	21	0	42	0	42
Blocks access to networked sensor information	0	5	3	5	3	0	0	8	8
Leverages jamming against fraction of time transmitting versus sensing	4	2	4	5	5	0	0	10	10
Induces receiver errors as if from a primary device	4	2	2	4	4	0	0	8	8
Jams at spectrum handoff or initiation	18	12	15	30	15	15	9	21	45
Misbehaves forwarding information between networked CR	0	0	14	7	7	4	4	6	14
Total	105	76	89	148	122	88	100	82	

Table 5 also provides a sum of risk over all 12 attacks. This is somewhat of an oversimplified view since the sum depends on our choice of 12 attacks and the assumptions about the 18 different design combinations. With these caveats we observe from the bottom row the relative vulnerabilities along each dimension. Along the network architecture dimension the centralized cooperative is the least vulnerable while the non-cooperative is the most vulnerable. This follows because the non-cooperative scheme does not benefit from the inherent redundancy provided in the cooperative architecture. The distributed cooperative suffers mainly from its vulnerability to misbehaving nodes. Along the spectrum access method dimension, overlay is more vulnerable than underlay. This is mainly because the underlay scheme is more robust to jamming at spectrum handoff and attackers that attempt to insert false policies. Along the spectrum awareness method dimension, each method has its specific weaknesses and strengths. Geolocate and access a database is the most vulnerable due to its dependence on the easily-jammed GPS. The other two methods pose similar overall risk. Detection and sensing is robust against most policy manipulation since we assume that it rarely accesses policy information, however it is vulnerable to several significant sensing specific attacks. Beacons are immune to the sensing and GPS attacks. However it is generally more vulnerable in the other attacks, mainly because the beacon itself is an attack target.

Table 5 can be somewhat misleading since it aggregates all architectures along a single dimension. The relative benefits of individual combinations are lost. Table 6 provides a sum over attacks for each of the 18 CR design combinations and is derived by summing each of the 18 columns in Table 4. In Table 6, the non-cooperative architecture is always more vulnerable than its equivalent

Table 6 Multi-dimensional risk values

	Overlay			Underlay		
	Beacon	Geolocate database	Detection sensing	Beacon	Geolocate database	Detection sensing
Non-cooperative	19	20	18	16	18	14
Centralized cooperative	12	17	13	10	15	9
Distributed cooperative	17	16	16	14	14	12

cooperative architecture in the same column. This suggests that from a security perspective the non-cooperative architecture should be avoided despite its simpler deployment. Similarly, the overlay network access method on the left is always more vulnerable than the equivalent underlay method on the right. This suggests that from a security perspective, the underlay method should be used despite its extra complexity. Comparing the spectrum sensing methods did not produce a clearly less vulnerable method. The four least vulnerable designs are (using obvious notation); (CC, DS, U); (CC, B, U), (CC, B, O), and (DC, DS, U). Interestingly, this set consists of (CC, DS, U) and 3 variations that vary one parameter along each of the three dimensions. We also note that two of the least vulnerable designs use the (CC, B) combination chosen for the IEEE 802.22 standard [19].

We wish to emphasize the limitations of this analysis. First it depends on the set of attacks examined. While we chose them to be representative they do ignore attacks such as attacks on the integrity of remote policy databases. The analysis only considers the security vulnerabilities and does not include cost or flexibility. Finally, only DoS security attacks are reviewed in this paper and not attacks on privacy, or unauthorized access [21].

7 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 can induce large performance degradation for relatively little effort. 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 relative risks for each combination of CR network architecture, spectrum access method and spectrum awareness method. 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. Akyildiz IF, Lee WY, Vuran MC, Mohanty S (2006) NeXt generation dynamic spectrum access cognitive radio wireless networks: a survey. Comput Networks 50:2127–2159
2. Barbeau M (2005) WiMax/802.16 threat analysis. In: Proceedings of the 1st ACM international workshop on quality of service & security in wireless and mobile networks, Quebec, Canada
3. Brown TX (2005) An analysis of licensed channel avoidance strategies for unlicensed devices. In: Proceedings of IEEE DySPAN, Nov. 8–11
4. Brown TX, James JE, Sethi A (2006) Jamming and sensing of encrypted wireless ad hoc networks. In: Proceedings of the seventh ACM international symposium on mobile ad hoc networking and computing (MobiHoc), Florence, 22–25 May
5. Brown TX, Sethi A (2007) Potential cognitive radio denial of service attacks and remedies. In: Proceedings of the international symposium on advanced radio technologies 2007 (ISART 2007), Boulder, 26–28 Feb
6. Brown TX, Sethi A (2007) Potential cognitive radio denial-of-service vulnerabilities and protection countermeasures: a multi-dimensional analysis and assessment. In: Proceedings of the second international conference on cognitive radio oriented wireless networks and communications 2007 (CrownCom 2007), Orlando, 31 July–3 Aug
7. Brown TX, Sicker D (2007) Can cognitive radio support broadband wireless access? In: Proceedings of the IEEE DySPAN, April 17–20
8. Chapin JM, Lehr WH (2007) Time-limited leases for innovative radios. In Proceedings of the IEEE DySPAN, April 17–20
9. DARPA XG Working Group (2003) "The XG vision," Request for comments, version 1.0, Prepared by BBN Technologies, Cambridge, Mass., USA. July
10. DARPA XG Working Group (2003) "The XG architectural framework," Request for comments, version 1.0, Prepared by BBN Technologies, Cambridge, Mass., USA. July
11. End-to-End Reconfigurability (E2R) Phase II website (e2r2.motlabs.com)
12. ETSI (2003) Telecommunications and internet protocol harmonization over networks (TIPHON) release 4; part I: threat analysis. Technical specification ETSI TS 102 165-1 V4.1.1
13. FCC ET Docket No. 02-135, "Spectrum policy task force report," Nov. 2002. (http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-228542A1.pdf)
14. 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)
15. 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
16. Heusse M, Rousseau F, Berger-Sabbatel G, Duda A (2003) Performance anomaly of 802.11b. In: Proceedings of the INFOCOM 2003, v. 2, 30 Mar.–3 Apr. 2003, pp 836–843
17. Householder A, Manion A, Pesante L, Weaver GM (2001) Managing the threat of denial-of-service attacks. CERT Coordination Center, v10.0, October
18. Hubaux J, Buttyán L, Čapkun S (2001) The quest for security in mobile ad hoc networks. In: Proceedings of the ACM symposium on mobile ad hoc networking and computing (MobiHOC), New York, NY
19. IEEE 802.22 Working Group on Wireless Regional Area Networks, (http://www.ieee802.org/22/)

20. Ma L, Han X, Shen C-C (2005) Dynamic open spectrum sharing MAC protocol for wireless ad hoc network. In: Proceedings of the IEEE DySpan, Nov. 8–11, pp 203–213
21. Mathur CN, Subbalakshmi KP (2007) Security issues in cognitive radio networks. In: Cognitive networks: towards self-aware networks, July
22. Menon R, Buehrer RM, Reed JH (2005) Outage probability based comparison of underlay and overlay spectrum sharing techniques. In: Proceedings of the IEEE DySpan, Nov. 8–11, pp 101–109
23. Newsome J, Shi E, Song D, Perrig A (2004) The Sybil attack in sensor networks: analysis & defenses. In: Proceedings of third international symposium on information processing in sensor networks, IPSN, 26–27 Apr., pp 259–268
24. Pilosof S, Ramjee R, Raz D, Shavitt Y, Sinha P (2003) Understanding TCP fairness over a wireless LAN. In: Proceedings of the INFOCOM 2003 v. 2, 30 Mar.–3 Apr., pp 863–872
25. Sankaranarayanan S, Papadimitratos P, Mishra A, Hershey S (2005) A bandwidth sharing approach to improve licensed spectrum utilization. In Proceedings of the IEEE DySPAN, Nov. 8–11, pp 279–288
26. Shirey R (2000) RFC 2828: internet security glossary. IETF, May
27. Stallings W (2006) Network security essentials: applications and standards, 3rd edn. Prentice Hall, Upper Saddle River, NJ, p 432
28. Ståhlberg M (2000) Radio jamming attacks against two popular mobile networks. In: Lipmaa H, Pehu-Lehtonen H, (eds) Helsinki University of Technology, Helsinki
29. Volpe JA (2001) Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System. Final report for the National Transportation Systems Center, US Department of Transportation, Aug. 29 (http://www.navcen.uscg.gov/gps/geninfo/pressrelease.htm)
30. Wenyuan X, Wade T, Yanyong Z, Timothy W (2005) The feasibility of launching and detecting jamming attacks in wireless networks. In: Proceedings of ACM international symposium on mobile ad hoc networking and computing, Illinois, pp 46–57

Timothy X. Brown received his B.S. in physics from Pennsylvania State University and his Ph.D. in electrical engineering from California Institute of Technology in 1990 when he joined the Jet Propulsion Lab. In 1992 he joined Bell Communications Research. Since 1995 he has had a joint appointment with the Department of Electrical and Computer Engineering and the Interdisciplinary Telecommunications Program at the University of Colorado, Boulder. He is currently an Associate Professor. His research interests include adaptive network control, wireless communications systems, and spectrum policy. He is a recipient of the NSF CAREER Award. In 2003 he was chosen the Global Wireless Education Consortium's (GWEC) wireless educator of the year.

Amita Sethi received her B. Tech degree from Mysore University, India in 1999. From January 2000 to September 2005, she has worked in the telecommunications software industry with Aricent Technologies (formerly, Flextronics Software Systems). Since January 2006, she is a Masters student at the University of Colorado, Boulder and is a research assistant in Professor Timothy Brown's wireless networking lab. Her research interests include security in cognitive radio networks and wireless ad-hoc networks.