TRIESTE: A Trusted Radio Infrastructure for Enforcing SpecTrum Etiquettes

Wenyuan Xu Pandurang Kamat Wade Trappe
Wireless Information Network Laboratory (WINLAB), Rutgers University
Piscataway, NJ 08854

***Abstract*— There has been considerable effort directed at developing "cognitive radio" (CR) platforms, which will expose the lower-layers of the protocol stack to researchers, developers and the "public". In spite of the great potential of such a radio platform, such "public" development threatens the success of these platforms: the proliferation of such wireless platforms, plus the open-source nature of their supporting software, is powerful but also dangerous. It is easily conceivable that inexpensive and widely available cognitive radios could become an ideal platform for abuse since the lowest layers of the wireless protocol stack are accessible to programmers. In order to regulate the future radio environment, this paper presents a framework, known as TRIESTE (Trusted Radio Infrastructure for Enforcing SpecTrum Etiquettes), which can ensure that radio devices are only able to access/use the spectrum in a manner that conforms to their privileges. In TRIESTE, two levels of etiquette enforcement mechanisms are employed. The first is an on-board mechanism that ensures trustworthy radio operation by restricting any potential violation operation from accessing the radio through a secure component located in each CR. External to individual cognitive radios, an infrastructure consisting of spectrum sensors monitors the radio environment, and reports measurements to spectrum police agents that punish CRs if violations are detected.**

I. Introduction

There has been considerable effort directed at developing "cognitive radio" (CR) platforms, which will expose the lower-layers of the protocol stack to researchers and developers [1]. This initiative is supported by two separate technical efforts: first, is a wealth of research devoted to uncovering the gains that are possible by letting the lower protocol layers become programmable and adaptable; second, are the recent advances in programmable integrated circuits that have significantly increased the amount of computation that can be done without requiring specialized hardware/firmware components. By being able to scan the available spectrum, select from a wide range of operating frequencies, adjust modulation waveforms, and perform adaptive resource allocation– all of these in real-time– these new "cognitive" radios will be able to adapt to a wide variety of radio interference conditions and adaptively select the most efficient communication mechanisms.

While there is great potential for such a radio platform, some caution regarding their ubiquitous use in wireless systems is warranted since their deployment will not be limited to the laboratory. Already, the GnuRadio platform [2] is available for general use, and supporting this platform is an open-source software effort to develop GnuRadio "blocks" [3]– software modules capable of conducting a broad range of functions associated with the reception/transmission of radio signals. Other CR platforms, such as the Xilinx-based Rice platform or the WINLAB-GaTech-Lucent cognitive radio platform, will also reach a large consumer base with similar open-source efforts supporting lower-layer protocols.

Such "public" development also threatens the success of these platforms: the proliferation of such wireless platforms, plus the open-source nature of their supporting software, is empowering but also dangerous. *It is easily conceivable that inexpensive and widely available cognitive radios could become an ideal platform for abuse since the lowest layers of the wireless protocol stack are accessible to programmers.* Thus, the gains promised by adaptive resource allocation schemes and good spectrum etiquette policies can be negated if cognitive radio devices can be reprogrammed to violate or bypass locally fair spectrum policies either maliciously or inadvertently. If fail-safe mechanisms are not employed, individual devices could use the wireless medium to their advantage at the expense of the greater good. It is therefore essential that these software radios have methods to ensure that the radio device and the implementations of their lower layer protocols are trustworthy, and that all cognitive radios are held accountable for not following locally acceptable spectrum etiquette.

In order to regulate this future radio environment, this paper presents a framework, known as TRIESTE (Trusted Radio Infrastructures for Enforcing SpecTrum Etiquettes), which will guarantee that a coalition of autonomous cognitive radios, each programmable and running its own suite of spectrum etiquette protocols, behaves according to acceptable communal policies. This paper will not study the design of optimal spectrum etiquette protocols, but instead will examine how to formalize the spectrum policies, and discuss mechanisms to enforce these policies.

We begin the paper in Section II by providing an overview of our proposed radio etiquette enforcement framework, TRIESTE. We then turn to focus on formalizing spectrum access control policies in Section III. In order to enforce spectrum access control policies, in Section IV and Section V, we introduce two complementary mechanisms. The first mechanism ensures trustworthy radio operation by restricting any violation attempt from accessing the radio through a secure on-board component. In the second mechanism, an external Distributed

The authors may be reached at {wenyuan, pkamat, trappe}@winlab.rutgers.edu.

Spectrum Authority (DSA) observes the radio environment, and will punish CRs if violations are detected. Finally, we wrap up the paper by providing concluding remarks in Section VI.

II. Architecture Overview

In this section, we focus on our framework, known as TRIESTE (Trusted Radio Infrastructure for Enforcing Spec-Trum Etiquettes), which is targeted at a generic cognitive radio scenario, as depicted in Figure 1. TRIESTE provides assurance regarding the operation of a cognitive radio. Towards this objective, TRIESTE is composed of three tiers of key actors: the lowest tier is the layer of general purpose cognitive radios, above that is the Distributed Spectrum Authority tier (DSA-tier), also known as spectrum police agents, and the top tier consists of the Spectrum Law Makers, e.g. FCC. For the purpose of discussion, we shall consider the Spectrum Law Makers as being secure (e.g. located at some secure central location), usually far away from the local radio environment. The cognitive radios and distributed spectrum authorities, such as spectrum police agents and auxiliary spectrum sensors are scattered in the local radio environment.

A. Spectrum Law Makers

The Spectrum Law Makers regulate the radio spectrum at the highest level; they define the laws, which will restrict the spectrum etiquette policies that are programmed by CR users. We note here the difference between laws and spectrum etiquette policies. The laws are made by the Spectrum Law Makers, which regulate how the radio spectrum should be accessed in general and serve as the guideline for spectrum etiquette policies. The spectrum etiquette policies are defined by individual cognitive radio users, or spectrum owners, and they have to obey the spectrum laws. For example, a law can mandate that an entity should not leak energy outside the spectrum it has negotiated; or, if one's spectrum sensor detects sufficient energy in a channel, then one cannot invade that channel unless the energy in that channel is caused by background interference. A sample spectrum etiquette policy

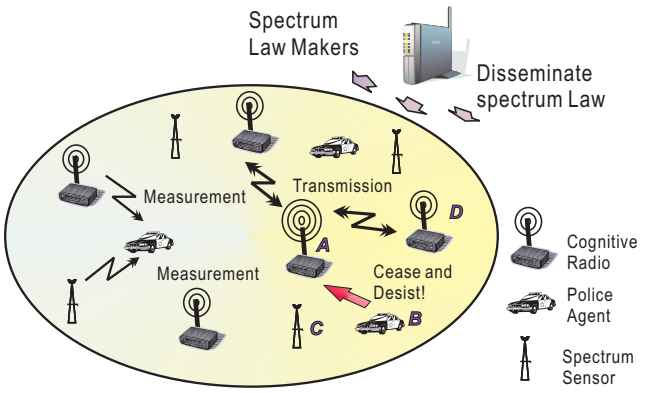

Fig. 1. TRIESTE Architecture is composed of multiple tiers of actors: lowest tier are the cognitive radios, above that will be the DSA-tier (Distributed Spectrum Authority), and above that are the Spectrum Law Makers (aka. FCC).

can be that an entity can continuously access an allocated spectrum band during a specified time period.

Besides spectrum access laws, the law makers should also make the punishment rules that are applied if any individual violates the spectrum laws. For example, if a radio device accesses spectrum that is not assigned to it, then it's privilege on another portion of the spectrum could be suspended for certain amount of time.

After defining the spectrum laws, they need to be disseminated to the distributed spectrum authorities and cognitive radios as well. To minimize the overhead of spectrum policy dissemination, it is desirable to automate the law dissemination procedure. Thus, the law should be defined in a format that is understandable by the radio device itself without human intervention. It is therefore necessary to have a formal language describing the laws that can restrict the possible spectrum etiquette policies programmed by CR users. Laws can be defined using an ontological representation (e.g. [4]), and similar representations can be used to specify access control policies (depending on the scenario, ranging from mandatory access control to contextual access control [5]). The DARPA XG effort has made initial headway into providing a language for specifying such laws/policies. We would recommend not intend to reinvent the wheel, but rather would suggest using the XG-OWL language to define upper/lower bounds on acceptable behaviors. By doing so, these acceptable behaviors will be passed down by the Spectrum Law Makers to the spectrum police agents as well as to the spectrum users (the cognitive radios) automatically.

We will discuss the detailed usage of XG-OWL, and how it may be used to specify limitations on acceptable actions over spectrum access as well as the punishment needed to control violations in Section III.

B. Generic Cognitive Radio and its On-board Enforcement

The Spectrum Law Makers give general guidelines on how the spectrum should be accessed. However, these laws need to be supported by enforcement mechanisms local to the cognitive radios. It is therefore necessary to have an on-board Trusted Computing base/module (TRIESTE-TCB) in each cognitive radio that enforces the spectrum laws and etiquettes.

The TRIESTE-TCB, as depicted in Figure 2, includes all the hardware and software in the cognitive radio that enforces universal laws and etiquette policies passed down by the Spectrum Law Makers. The TRIESTE-TCB can be thought of as the controlled gate that users have to go through to access the radio. In TRIESTE, typically, before the user can transmit information over a certain radio spectrum band, the user/process has to send a spectrum access request, which includes information about the target radio frequency band, the spectrum etiquette the user will follow, the transmission power, transmission duration, etc. to the packet processor. Here, we note that we shall abuse terminology and, for simplicity, collectively refer to the packet processor as an entity consisting of multiple processors handling packets, such as the Network Processor, the CR Policy Processor etc. The packet processor

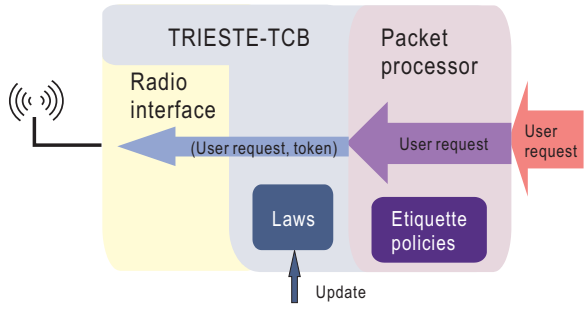

Fig. 2. The architecture of the Cognitive Radio with on-board TRIESTE-TCB.

shapes the user radio access request according to the spectrum etiquette policies programmed by the user or spectrum owner, then passes the modified user request to TRIESTE-TCB. The TRIESTE-TCB in turn will validate the request against the laws available to it and will allow the request to go through only if it does not violate any of those laws.

C. Distributed Spectrum Authority

On-board law enforcement can, in general, secure the spectrum access. However, at some point we must confront the possibility of a truly greedy or even malicious users who can circumvent the safeguards put in place to help the "honest" cognitive radios avoid breaches of etiquette. In order to cope with these more serious threats, we propose to enforce spectrum law through means external to the cognitive radio itself. This is represented by the second tier in TRIESTE Architecture shown in Figure 1, which we refer to as the Distributed Spectrum Authority (DSA).

The DSA or the "spectrum police" monitor the local radio environment by collecting geographically distributed radio measurements from the population of cognitive radios as well as auxiliary spectrum sensors. Of course, since some measurements can be supplied by potentially greedy/rogue users, there is a risk that such biased data might influence the actions of the DSA. One interesting research direction involves specifying techniques to filter out inaccurate data, reliably access an interference environment, and detect violations by comparing with "spectrum laws" (e.g. issued by FCC).

As a violation is detected by the spectrum police, corresponding local punishments, either condoned or coordinated by the DSA, are enacted. To enforce proper spectrum law, the distributed spectrum authority could shutdown offending CRs via an authenticated kill-switch located on each cognitive radio. In some extreme cases, where the kill-switch located on the cognitive radio is disabled by a malicious user, a further level of enforcement can be conducted by utilizing RF-localization techniques and seizing rogue transmitters.

As an example, consider the cognitive radio A, depicted in Figure 1, which keeps transmitting a signal in a spectrum band that hasn't be assigned to it. The spectrum police agent B collects the radio measurements from its neighbor cognitive radios, such as D and the spectrum sensors, e.g. C, located near it. The police agent then compares the radio activity with the Spectrum Laws. Based on its assessment, it detects that cognitive radio A should not have access to the spectrum, and it sends out a "cease and desist" command to the kill-switch on cognitive radio A. Thus, the cognitive radio A stops its radio as commanded. The effect of a cease and desist command could range from temporarily shutting down radio functionality (until a trusted punishment timer expires), or could involve completely disabling the CR until the user brings the CR into a central authority for reinstatement.

III. Spectrum Law/policy Formalism

In the TRIESTE framework, cognitive radios must adhere to the Spectrum Laws and spectrum etiquette policies that apply to their operation. Traditionally, laws/policies are published in a human readable form, and interpreted by humans. Those laws pertaining to RF devices are often hard-coded into radio devices. This is an inefficient, inflexible and non-scalable solution that does not facilitate the broadest range of spectrum usage and coexistence. Future cognitive radios should be able to adapt to new laws/policies dynamically, as laws/policies tend to change over time. It is therefore desirable to publish laws/policies in a well-defined language, and let cognitive radios interpret them without human intervention. Although laws and spectrum etiquette policies are produced by different entities, cognitive radios have to adhere to both of them. Therefore, they should be expressed in the same format to facilitate integration, consistency checking and conflict resolution. For simplicity, in this section, we use the term spectrum policies to refer to both spectrum laws and spectrum etiquette policies.

A. XGPL

We propose to use XGPL (XG Policy Language) [4] to express spectrum policies formally. XGPL is part of the XG (neXt Generation Communications) research program, which aims to let radios utilize available spectrum intelligently and dynamically based on the knowledge of actual conditions and spectrum policies. In particular, the XG project chose OWL(Web Ontology Language) as its XG Policy Language for several reasons. First of all, OWL provides the structure and richness needed to express policies. Secondly, general theorem proving/reasoning engines for deductive interference are already available. Finally, OWL is an efficient language for describing data, and passing data around different systems.

OWL is originally designed for processing information on the Web and is designed to be interpreted by computers. It is written in XML(Extensible Markup language). We note that OWL is not another programming language, but is a structured way to build representations for information and policies for machine understanding. For example, the OWL expression of magnitude is 10 is the following:

```
<xgparam:magnitude>
    <xsd:integer rdf:value="10" />
</xgparam:magnitude>
```

The paragraph above defines a property "magnitude" in the name space "xgparam". The value of the property magnitude is 10, the type of the value is integer, which is defined in

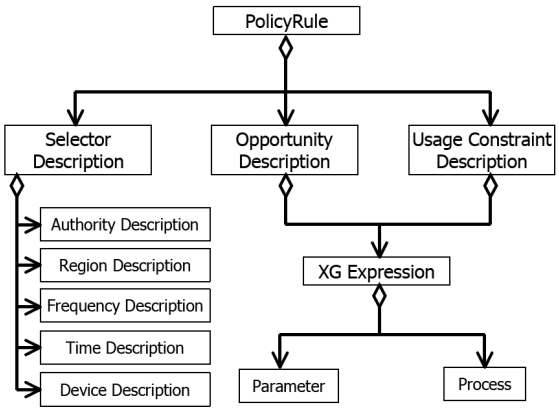

Fig. 3. Structure of Policy Facts as described from [4].

namespace rdf. More detailed and precise exposition on OWL can be found in [6].

In this paper, instead of expressing policies in the original OWL language, we use the shorthand notion, which is also used in [4]. The shorthand notation yields representations equivalent to OWL representations. For example, we describe the previous "magnitude is 10" in the following way:

```
(magnitude 10)
```

Detailed mapping from OWL to shorthand notion can be found in [4].

B. XGPL Elements

XGPL language is a declarative language based on facts and rules instead of a procedural language. The spectrum policies defined are a set of declarations instead of many layers of "`if-then-else`" clauses. For example, the set of spectrum policies (in English) will look like the following (we will present the XGPL expression later in this section):

- Transmissions shall be contained within $3.6GHz$ to $3.7GHz$
- The peak power spectral density shall not be more than $1nW/Hz$
- etc. ...

A spectrum policy rule is composed of three facts: a selector description, an opportunity description and a usage constraint description, as shown in Figure 3.

The first part in a spectrum policy rule is a selector description, which is used to filter policy rules to the sub set of rules that may apply to a given situation. The selector description contains one or more facts that describes the frequency, time and region the policy covers, the authority that define the policy, and the radio device to which the policy rule applies. For example, a selector description may include filters such as "applies to operation in U.S.A" or "applies to operations in the $3.6GHz$ to $3.7GHz$ bands".

The second part in a policy rule is an opportunity description, which is used to evaluate whether the transmission request is valid or not based on whether or not a given environment and device state matches the opportunity description in the filtered subset rules. For example, the opportunity description can be "if a beacon is heard at $823MHz$", or "peak received power is less than $-80dBm$".

A valid opportunity indicates transmission that conforms to the usage constraint description is permitted. Usage constraint description constrains the radio behavior, such as "transmit with a maximum power of $-10dBm$" or "maximum continuous on-time must be 1 second and the minimum off-time must be $100msec$".

We envision that, usually, a spectrum policy rule is first defined in XGPL, then each element (a selector description, an opportunity description and a usage constraint description) is defined in a format similar to the format used to specify policy rules.

C. Example

Now let's look at an example on using XGPL to formalize spectrum policies. Suppose we have the following spectrum policies in English:

- This policies apply to CR devices that are capable of operating in the $3.6 - 3.7GHz$ band
- Transmissions shall be contained within $3.6GHz$ to $3.7GHz$
- The peak power spectral density shall not be more than $1nW/Hz$

Translating the policies above into XGPL, we get the eight facts `L1-L8:` as listed in Table I. `L1` and `L2` define two policy rules with the identifier `(id)` `P1` and `P2`. Each policy rule contains a selector, opportunity and usage description, which are defined as facts as well. For example, policy rule `P1` points to selector description `S`, defined in `L3`, opportunity description `AnyOpp`, whose definition is omitted here, and usage description `U_Band`, defined in `L4`. Here, opportunity description `AnyOpp` means any opportunity, or any situation is applied. `L3` defines selector S that includes three descriptions `FreqDesc, RegnDesc`, and `TimeDesc`. The frequency description, `FreqDesc` is defined in `L6-L7` as a frequency band ranging from $3.6 - 3.7GHz$. To emit a signal, the emission operation has to conform to the usage description defined in `L4-L5`, e.g. the emission band has to be within the range of `XGCBand`, and the peak power spectrum density (PSD) has to be less than $1nW/Hz$.

Assume we have a user request such as "want to transmit at $3.6GHz$ with the PSD $2nW/Hz$". First, policy filtering is performed, the resulting policy set will include `P1` and `P2`. Further, since there are no specific opportunity descriptions, no policy rules are filtered out by opportunity description matching. Next, the usage descriptions are checked. The user request can pass the band usage check, however, it will fail over the PSD usage description, because $2nW/Hz > 1nW/Hz$. Therefore, this user request will not pass the policy checking in the end, and no access permission will be granted.

D. Punishment Policies

The spectrum law includes both Spectrum Access laws, and punishment laws. In the previous subsection, we discussed how to express spectrum access rules using XGPL. In the

Index	XGPL description
L1	(PolicyRule (id P1) (SelDesc S) (OppDesc AnyOpp) (UseDesc U_Band))
L2	(PolicyRule (id P2) (SelDesc S) (OppDesc AnyOpp) (UseDesc U_PSD1))
L3	(SelDesc (id S) (FreqDesc XGConopsBand) (RegnDesc US) (TimeDesc Forever))
L4	(UseDesc (id U_Band) (xgx "(within Emission.Band XGCBand)"))
L5	(UseDesc (id U_PDS1) (xgx "($\leq$ Emission.PeakPSD PeakPSD1)"))
L6	(FreqDesc (id XGConopsBand) (FrequencyRange SGCBand)
L7	(FrequencyRange (id XGCBand) (minValue 3.6) (maxvalue 3.7) (unit GHz))
L8	(PSD (id PeakPSD1) (magnitude 1.0) (unit nWperHz))

TABLE I
POLICIES EXPRESSED IN XGPL

original XG project, XGPL is designed to describe spectrum access control. XGPL was not used to specify any form of punishment for spectrum abuse. In particular, the underlying idea of the XG project is that the regulatory policy does not tell the radio what to do, it only defines what constitutes authorized use of the spectrum. Punishment, however, tells the radio what should be done once violation is conducted.

We believe that it is necessary to define punishment rules as part of the spectrum laws, since punishment can serve as a precaution against potential spectrum violation as well. Although it might be a challenge, XGPL can be extended to define punishment rules. One way to define punishment is to add one more description, punishment description into the policy rules as shown below:

```
(PolicyRule (id Policy_name)
    (SelDesc S)
    (OppDesc SomeOpp)
    (UseDesc SomeUseDesc)
    (PunDesc SomeAction)
```

One possible way to perform the punishment is as follows. If the punishment rule is selected and activated, then new punishing rules with certain expiration period will be generated based on the level and type of punishment, and inserted into the existing spectrum polices for specified amount of time. For example, the newly generated punishing spectrum

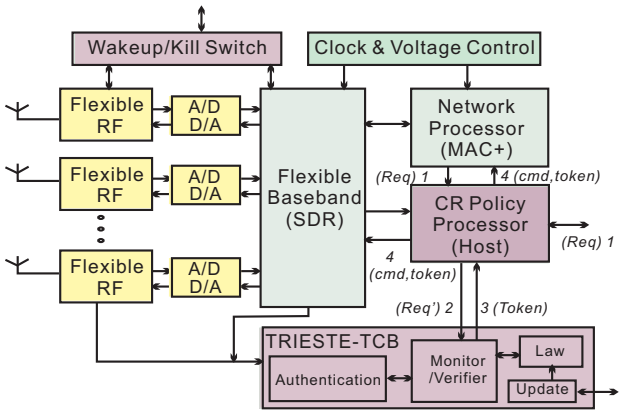

Fig. 4. A generic SDR/CR platform involving RF processors, baseband processor, network processor, and the cognitive radio processor. Note that TRIESTE regulates via a TCB component and an externally-accessible authenticated kill-switch.

access rules could be that the radio device cannot access to band $3.6 - 3.7GHz$ for two hours. Of course, precedence mechanisms are needed in order to resolve conflict. Detailed techniques for defining punishment and precedence require further investigation.

We note that punishment rules don't need to be implemented in the on-board enforcement, i.e. TCB, since the major goal of the on-board TCB is to prevent illegal spectrum access. Rather, punishment rules are more of an external mechanism that the spectrum police agent uses to take action against violations already performed, and after damage has already been done to other legal spectrum users.

IV. ON-BOARD ENFORCEMENT

The cognitive radio, is a programmable wireless platform that will support a wide range of radio network scenarios from autonomous agile radios to those that use higher layer protocols to share spectrum. A typical high-level architecture for a CR is shown in Figure 4. It consists of Flexible RF units, a baseband processor, a network processor and a cognitive radio policy processor (which also functions as the host). Besides those components, we have added a logical component, the TRIESTE-TCB, to enforce the spectrum laws. Here we want to point out that the law/policy enforcement activities are likely to be performed at several functional places within the CR, as law/policy enforcement is potentially related to every network protocol that will access the spectrum. Although we show the TRIESTE-TCB in one monolithic block, in implementation, the functions of the TRIESTE-TCB will be located in firmware in different processors. We now review some of the basic components in a CR.

RF Processor: The flexible RF front-end serves as both the receiver and the transmitter. It is the bridge between radio channel and hardware. The RF front, connected to antennas, can receive the signal from the channel and feed the analog signal to the A/D converters. Alternatively, the RF front-end takes the signal from the D/A converter and transmits the signal to the air. Typically, the RF front-end consists of a number of reconfigurable transceivers, each of which provides a narrow band channel (5-20 MHz) tunable across a certain

spectrum range. The RF transceivers will be programmable in terms of carrier frequency, bandwidth and transmit power– all parameters that can be tuned by spectrum etiquette algorithms (or by an adversary).

A/D and D/A converters: A set of analog to digital (A/D) and digital to analog (D/A) converters interface the RF front-ends to the baseband subsystem.

Baseband Processor: In addition to supporting regular wireless baseband functions, this processor will be used to support spectrum scanning and analysis.

Network Processor: The network processor performs MAC and higher layer packet processing. In addition to controlling access to the wireless physical layer, it functions as a router between the various radio links and can provide a fast ethernet port used to connect to the wired infrastructure. When receiving data, physical layer frames flow from the baseband processor, the header is separated from the payload and they are stored in different buffers respectively for higher layer to use.

CR Policy Processor: The CR policy processor is programmable by the user and adjusts other processors by downloading code to these processors or performs necessary reconfiguration of the processors. It uses the other processors to set up wireless links and implements collaborative networking sessions with other nodes. These adjustments will be regulated using a trusted computing module that enforces spectrum laws and etiquette policies specified using the formalisms described in Section III. The CR policy processor works with a set of spectrum etiquette policies that are defined by the spectrum owner. We envisage at least two kinds of scenarios that may arise here, namely licensed and unlicensed spectrum usage. In the licensed spectrum case the user has no administrative control over the spectrum policies that can be used. These policies would be specified by the spectrum licensee and loaded into the CR processor along with a signed certificate of the policy. The CR processor verifies the signature on policies that deal with licensed spectrum. The licensee's public keys can be loaded into the CR processor through the secure software download mechanism already in place for the rest of the CR system. For the unlicensed spectrum the user can have a higher level of control in defining the spectrum etiquette policy. This would be allowed in order to facilitate research and experimentation with the CR platform and building of new applications.

TRIESTE-TCB: The TRIESTE-TCB can be thought as the controlled gate that users have to go through to access radio. The basic structure of TRIESTE-TCB consists of a generic *Controller* which can interpret and enforce any well-formed *Law*. As we pointed out earlier, the TRIESTE-TCB is a virtual block, and the real functions of the TRIESTE-TCB will be located in hardware or software on different components of the CR.

Conceptually, the TRIESTE-TCB evaluates the access request along with the user's credentials and checks it against the spectrum laws. If the request and credential combination is valid in the context of the current spectrum laws, then the TCB issues a privilege token for that request, as shown in Figure 4 step 3. The privilege token is a tuple consisting of $\langle$ spectrum-access-details, timestamp, and a signed hash of (spectrum-access-details $\|$timestamp)$\rangle$. The spectrum-access-details contain the radio frequency, duration, and spectrum access limitation granted. If the user's credentials don't permit the privilege level of the request or if the combination somehow violates some spectrum law, then the TRIESTE-TCB could either try to find a permissible modification of the request that is in compliance with the spectrum laws or reject the request if such a modification is not feasible. One interesting feature here is that the user's credentials may change over time and each request is evaluated in the context of the credentials presented to it. For example, a user with emergency-responder credentials would have a higher privilege to spectrum access during an emergency situation as opposed to during a non-emergency situation.

To understand how the TCB works, let us go through an example. Initially, after the user request (req) comes in, as shown in Figure 4 step 1, from the Network Processor or directly into CR Policy Processor, the CR Policy Processor shapes the request into req' as described above. Then, req' is sent into TRIESTE-TCB, the Monitor/Verifier module in TRIESTE-TCB evaluates the access request along with the user's credentials and determines whether the privilege token can be issued. In order to optimize the performance, a batch of user operations should be submitted to Monitor/Verifier together. However, tokens should be associated with an atomic spectrum operation, which is the smallest operational unit that is processed by Flexible Baseband(SDR). In this way, the Flexible Baseband can authenticate the operations one by one.

Inside TRIESTE-TCB, the Monitor/Verifier module will also monitor the on-board radio activity, observe the radio environment, and check any potential violation by comparing against "spectrum laws". If, in the very rare case, the user doesn't follow the etiquette it claims to obey and thus violates the spectrum law, the TRIESTE-TCB will stop the radio operation and revoke the user's token/privilege. In this case, the request will not have the requisite credentials, and will not be acted upon by the rest of the CR.

The Laws that the TCB works with, as well as basic security parameters (i.e. cryptographic keys), should be stored in a secure storage container, to protect against tampering. Additionally, after the token has been issued to the user, the association relationship between user request and the token should not be altered while "(user request, token)" pair is being passed among cognitive radio components. To achieve integrity, message authentication codes can be used. Alternatively, we can design the cognitive radio in such a way that, after the creation of "(user request, token)", the pair travel through the components via trusted paths. Thus, since the data pair cannot be intercepted on the way, the content of the user request also cannot be changed. Additionally, as the spectrum laws will evolve over time, it is desirable to make the law's stored on the CR upgradable. The TCB will download a new law only if it is signed by the regulating authority, such as the FCC. As usual, the more flexible the system is, the more layers of security and associated cryptographic material, need to be integrated into the TCB.

Wakeup and Kill Switch: A "wakeup" module allows one

or more of the radio channels to respond to simple coded messages and bring the baseband processor out of a deep (low power) sleep. The "kill" module takes direct commands issued by spectrum police and stops the corresponding ongoing radio activities. The wakeup and kill switch module should be authenticated in order to assure that only valid wakeup/shutdown commands are acted upon.

V. External Infrastructure for Enforcing Radio Behaviors

The on-board enforcement mechanisms can prevent most of the spectrum violations. We must, however, confront the possibility of truly greedy or even malicious users who can circumvent the safeguards put in place to help the "honest" cognitive radio avoid breaches of etiquette. In this section we will explore a variety of research questions associated with enforcing spectrum etiquette through means external to the cognitive radio itself. The issues associated with externally regulating radio behavior range from whether a Distributed Spectrum Authority (DSA) can reliably assess an interference environment based on a collection of geographically distributed measurements (with some measurements being supplied by potentially greedy/rogue users), to local punishment inflicted by transceivers (condoned or coordinated by the DSA), to the ultimate enforcement by localizing and seizing rogue transmitters.

A. Discovering the Crime

Trusting that all cognitive radios will report accurate measures of the radio environment to an Authority (or report anything at all) seems foolhardy. Cognitive radios, especially those behaving badly, have a vested interest in disseminating misinformation (to the police) which furthers their own ends. Thus, in the interest of the greater good, some effort must be made to ascertain whether what is being reported accurately reflects the reality.

One way to provide accurate readings is to deploy a dedicated and secure network which can probe and measure the environment. Of course, such a network could prove costly to deploy in numbers great enough to guarantee accurate readings. However, without *some* completely trustworthy readings, the Authority can depend only on self-interested reports.

Therefore, one research problem is to understand the necessary density of a dedicated sensor/probe network. With some ability to directly "sniff" the environment, networks of trusted cognitive radios could be developed through reputation – with reports measured against authority sensor readings when available. Since sensors might not necessarily emit radio energy (a trusted sensor network should probably have reliable and secure "wired" links to the authority), it will be difficult for mobile transceivers to know whether they are in the vicinity of a sensor. Thus, reporting inaccurate readings of the environment would carry the risk of reputation damage.

B. Identification

The notion of sensing crime leads directly to the issue of identifiability. Conventional criminology employs a variety of techniques to identify individuals associated with a crime, with simple fingerprinting being one of the most recognizable methods.

MAC addresses are the conventional means for identifying network devices. However, it is well known that such identifiers are insecure and can be set to arbitrary values by a savvy programmer (and further evidenced by the numerous security threats possible because of MAC address spoofing!). When we examine the issue of MAC addresses for cognitive radios, several questions arise. First, is the basic question of what the purpose of a MAC address is. If cognitive radios are truly meant to interface between multiple radio technologies, which use different formats for specifying MAC addresses, then what is really needed is an alternative identifier for cognitive radios– a universal identifier of sorts. Further, given the generality and power of the programming interface supplied by cognitive radios, this universal identifier must be made to be unalterable, and hence this phase of packet formation must be controlled by the TCB.

Additionally, some some other means of identifying transmitters tied directly to hard-to-alter characteristics might be desirable to consider. In military applications, a technique called Specific Emitter Identification (SEI) [7] is used to distinguish friendly radars from those of the enemy. SEI uses electromagnetic signature from a radar transmitter to distinguish between different radars. Inspired by SEI, recent work, such as the work done at the National Institute for Standards and Technologies (NIST) [8], attempts to measure transmitter radiative signatures from common wireless cards based on unavoidable and random fabrication differences. They observe that there are quantifiable differences both in the time and frequency domain between different wireless cards (even from the same manufacturer). However, the variability of the quantifiable differences depending on the physical location and orientation of the transmitter within an environment (due to multipath profiles, orientation of the antennas and the like) make it harder to apply this technique in a field system. Further study must be conducted to design a reliable method for identifying individual cognitive radios.

Taking motivation for identification from conventional forensic fingerprinting, it might be desirable to embed hard-to-alter RF fingerprints within every transmission coming from a particular cognitive radio. This is similar techniques used by the watermarking community to track, detect, and identify manipulations to media content, and similar signal processing techniques may be employed. It is an open research problem to study whether it is possible to transmit such unique RF signatures via the main communication band without significantly distorting the information coding. An alternative would be to transmit the RF signature out-of-band. The injection of such RF signature would be beyond the control of the user and controlled exclusively by the TRIESTE-TCB.

C. Punishment

Assuming radios can be identified, the issue of curbing illicit behavior arises. One can imagine a variety of different methods. For radios whose protocol somehow evaded the law-based verification process, something as simple as notification

of an etiquette violation could be sufficient– for example, a cease-and-desist message or a message warning that continued spectrum violations will result in more serious enforcement procedures. Of course, this raises the question of how such warning messages could be interpreted by the software of a CR, and this might lead to situations where an adversarial CR programmer implements a protocol that "pushes the limit" by performing violations until it receives the warning message, in much the same way as automobile drivers continue to speed until they are issued their first ticket. A simpler and more direct approach to punishment might be to employ a a remote "kill switch" which renders the radio inoperative. Effort must be taken, however, to secure the remote "kill switch". One way to make sure that only authorized entities can turn on the "kill switch" is to authenticate every remotely issued kill command. This requires that cryptographic material, such as public key parameters associated with the spectrum police, needs to be maintained and updated in a secure storage location on the CR.

A different, less direct, approach to enforcing spectrum laws would be to employ social phenomena, such as reputation-based systems. For example, in a network of cognitive radios, the reputation of each cognitive radio can be scored in a distributed manner. Individual CRs will only participate in multi-party operations (such as the forwarding of a packet) if the communication comes from a reputable/trustworthy CR. Hence, should a CR wish to have its transmissions forwarded or acted upon, then it must build and maintain a positive reputation by not acting in ways contrary to the accepted spectrum etiquette or spectrum policies. Such a method of soft-enforcement raises many interesting research questions, not the least of which is how to manage the formation of reputation.

There is also the possibility of creating deliberate interference for rational rogue transmitters using a dedicated infrastructure and/or honest "deputy" users. Since one can assume that the intended receiver is within some range of the rogue transmitter, the problem becomes one of identifying rogue transmissions and jamming them with Authority transmitters within "earshot." One obvious research problem is how vigorously (in time, space and frequency) to pursue such actions since jamming is a purely punitive action which does not carry traffic. In particular, jamming as a form of punishment can have a broader cost associated with it, although we might jam adversarial nodes we could also jam legitimate nodes.

D. Localization and Seizure

In the case of adversarial transmitters that are not able to be controlled by various forms of punishment, it might be then necessary to localize the misbehaving device and seize it.

In order to localize the transmitter, there is a large body of works devoted to the radio localization for a variety of different wireless network scenarios [9]–[16]. However, it must be realized that an adversarial device will not wish to be localized, and consequently the adversarial device will attempt to make its localization as difficult as possible. For example, an adversary might not use an isotropic antenna pattern, thereby altering the manner in which it distributes RF energy in the environment, and as a result make inferring its location more difficult (many localization schemes employing signal strength measurements assume isotropic radiation). The recent work on secure localization can server as a starting point for building reliable localization functionality [16]–[18].

Following the localization procedure, a real-world agent must apprehend the transgressor and impound the CR. This raises many issues, the least of which is the cost associated with performing such mundane activities. Additionally, issues such as the duration of an impound must be specified, along with a means for the CR owner to reclaim its CR after the punitive period has transpired. All in all, the localization and seizure procedures might only be considered as a method of last resorts, when the violations are so extensive and severe that the violating CR must be removed. In such cases, the need to localize should be infrequent, and since the violation was severe it would thus not be necessary to return the CR to its owner.

VI. Concluding Remarks and Future Directions

As Cognitive Radios (CRs) become ubiquitous in the future there will be attempts to misuse the highly open and granular control provided to the radio interface. In this paper we proposed a framework TRIESTE to secure the future radio environment by ensuring that radio devices are only able to access/use the spectrum in a manner that conforms to their privileges. We have presented methods to formalize spectrum laws and etiquettes, and discussed two levels of etiquette enforcement mechanisms. One is an on-board trusted computing base/module (TCB) and the other is an infrastructure external to individual CR, consisting of spectrum sensors and spectrum police agents. The on-board TCB prevents potential violation operations from accessing the radio by comparing them with a set of predefined spectrum laws and etiquettes. The external infrastructure monitors the radio environment, punishes CRs if violations are detected, and in some extreme case, localizes and seizes the rogue transmitters.

Future work in this direction must parallel the development of cognitive radios, and in particular must further map out architectural and systems issues. There are several separate research directions that should be explored. First, would be to conduct a prototyping effort to evaluate the impact of using an initial TRIESTE-TCB on the performance of a CR. Such an effort would help map out the interplay between policies, their interpretations, and their enforcement using onboard mechanisms. Next, is to examine several different identification mechanisms that can facilitate the recognition of CRs. Introducing RF signatures into a CR's transmissions, much like watermarks are applied to media, is an interesting area for investigation, but the underlying issue of whether this can be done in a provably secure manner without disrupting information coding remains to be seen. Another direction involves developing an integrated localization and identification system, using spectrum sensor readings and the cooperation of neighboring CRs.

REFERENCES

[1] G. Maguire and J. Mitola, "Cognitive radio: Making PCS personal," *IEEE PCS Magazine*, vol. 6, no. 4, pp. 13–19, 1999.

[2] "Gnuradio homepage," http://www.gnu.org/software/gnuradio/.

[3] "How to write a signal processing block for gnuradio," http://www.gnu.org/software/gnuradio/doc/howto-write-a-block.html.

[4] "Xg policy language framework, version 1.0," XG Working Group Document, prepared by BBN Technologies, 2004.

[5] M. Bishop, *Computer Security: Art and Practice*, Addison Wesley, 2003.

[6] "http://www.w3.org/tr/owl-features/," .

[7] L. Langley, "Specific emitter identification (sei) and classical parameter fusion technology," *WESCON '93*, 1993.

[8] K.A. Remley, C.A. Grosvenor, R.T. Johnk, D.R. Novotny, P.D. Hale, M.D. McKinley, A. Karygiannis, and E. Antonakakis, "Electromagnetic Signatures of WLAN Cards and Network Security," in *5th IEEE International Symposium on Signal Processing and Information Technology (IEEE ISSPIT 2005)*, December 2005, Athens.

[9] K. Dogancay, "Emitter localization using clustering-based bearing association," *"IEEE Transactions on Aerospace and Electronic Systems"*, vol. 41, no. 2, pp. 525–536, April 2005.

[10] S.D. Coutts, "3-d emitter localization using inhomogeneous bistatic scattering," in *Proceedings ICASSP'99*, March 1999, pp. 1509–1512.

[11] J.J Jr. Caffery and G.L. Stuber, "Overview of radiolocation in cdma cellular systems," *"IEEE Communications Magazine"*, pp. 38–45, April 1998.

[12] A.J. Weiss, "On the accuracy of a cellular location system based on rss measurements," *"IEEE Transactions on Vehicular Technology"*, vol. 52, no. 6, pp. 1508–1518, November 2003.

[13] K. Langendoen and N. Reijers, "Distributed localization in wireless sensor networks: a quantitative comparison," *Comput. Networks*, vol. 43, no. 4, pp. 499–518, 2003.

[14] N. Priyantha, A. Chakraborty, and H. Balakrishnan, "The CRICKET location-support system," in *Proceedings of the 6th annual international conference on Mobile computing and networking (Mobicom 2000)*, 2000, pp. 32–43.

[15] D. Nicelescu and B. Nath, "Ad hoc positioning (APS) using AOA," in *Proceedings of IEEE Infocom 2003*, 2003, pp. 1734 – 1743.

[16] S. Capkun and J.P. Hubaux, "Secure positioning in sensor networks," Technical report EPFL/IC/200444, May 2004.

[17] Z. Li, W. Trappe, Y. Zhang, and B. Nath, "Robust statistical methods for securing wireless localization in sensor networks," in *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks (IPSN 2005)*, 2005, pp. 91–98.

[18] S. Capkun and J. P. Hubaux, "Secure positioning of wireless devices with application to sensor networks," in *Proceedints of IEEE INFOCOM 2005*, 2005.