Technical Document on Cognitive Radio Networks

GDD-06-20

GENI: Global Environment for Network Innovations

September 15, 2006

Status: Draft (Version 1.0)

Note to the reader: this document is a work in progress and continues to evolve rapidly. Certain aspects of the GENI architecture are not yet addressed at all, and, for those aspects that are addressed here, a number of unresolved issues are identified in the text. Further, due to the active development and editing process, some portions of the document may be logically inconsistent with others.

This document is prepared by the Wireless Working Group.

Editors:

Joe Evans, *U. Kansas*

Gary Minden, *U. Kansas*

Ed Knightly, *Rice*

1. Introduction

A driving feature of future network architectures will be the mobile user. Users increasingly will access information resources while on the move, whether when in a vehicle, attending a business meeting, or working in remote locations. Wireless technology is necessary to support the mobile user and adaptive and efficient use of radio spectrum is an important aspect of developing future network architectures. Cognitive Radios (CRs) integrate radio technology and networking technology to provide efficient use of radio spectrum, a natural resource, and advanced user services.

The cognitive radio wireless network is intended as an advanced technology integration environment with focus on building adaptive, spectrum-efficient systems with emerging programmable radios. The emerging cognitive radio scenario is of current interest to both policy makers and technologists because of the potential for order-of-magnitude gains in spectral efficiency and network performance. NSF and industry funded R&D projects aimed at developing cognitive radio platforms are currently in progress and are expected to lead to equipment that can be used for GENI in the 2007-08 timeframe. Protocol research to be supported with the planned experimental system includes discovery and self-organization, cross-layer protocols for PHY adaptation, cooperation and competition mechanisms, spectrum etiquette procedures, and cognitive radio hardware/software performance optimization.

Cognitive Radio Features

The idea of a cognitive radio extends the concepts of a hardware radio and a software defined radio (SDR) from a simple, single function device to a radio that senses and reacts to its operating environment. A Cognitive Radio incorporates multiple sources of information, determines its current operating settings, and collaborates with other cognitive radios in a wireless network. The promise of cognitive radios is improved use of spectrum resources, reduced engineering and planning time, and adaptation to current operating conditions. Some features of cognitive radios include:

Sensing the current radio frequency spectrum environment: This includes measuring which frequencies are being used, when are they used, estimating the location of transmitters and receivers, and determining signal modulation. Results from sensing the environment would be used to determine radio settings.

Policy and configuration databases: Policies specifying how the radio can be operated and physical limitations of radio operation can be included in the radio or accessed over the network. Policies might specify which frequencies can be used in which locations. Configuration databases would describe the operating characteristics of the physical radio. These databases would normally be used to constrain the operation of the radio to stay within regulatory or physical limits.

Self-configuration: Radios may be assembled from several modules. For example, a radio frequency front-end, a digital signal processor, and a control processor. Each module should be self-describing and the radio should automatically configure itself for operation from the available modules. Some might call this “plug-and-play.”

Mission-oriented configuration: Software defined radios can meet a wide set of operational requirements. Configuring a SDR to meet a given set of mission requirements is called mission-oriented configuration. Typical mission requirements might include operation within buildings, substantial capacity, operation over long distances, and operation while moving at high speed. Mission-oriented configuration involves selecting a set of radio software modules from a library of modules and connecting them into an operational radio.

Adaptive algorithms: During radio operation, the cognitive radio is sensing its environment, adhering to policy and configuration constraints, and negotiating with peers to best utilize the radio spectrum and meet user demands.

Distributed collaboration: Cognitive radios will exchange current information on their local environment, user demand, and radio performance between themselves on a regular bases. Radios will use their local information and peer information to determine their operating settings.

Security: Radios will join and leave wireless networks.

Cognitive Radio Research Questions

The capabilities outlined above for cognitive radios, and the questions they generate, are just forming in the wireless research community. In a truly integrated network, once one node is mobile, the entire network needs to deal with mobility. Routing is one of the first functions that must address mobility. Resource management under the constraints of mobile users is another open issue. A sample of research questions include:

1. How does a richly connected mobile wireless network affect the architecture, design, and implementation of a global internet? How do local resource allocation algorithms, say the allocation of RF spectrum resources, interact with fixed infrastructure resource allocation?
2. How does one reliably sense the radio spectrum environment? How does one detect weak signals with constrained processing capacity? How does one use measurements to set the operating parameters of the radio? How does one reliably exchange measurement information with peers?
3. How does one express regulatory and operational policies? How are policies securely updated? What method does one use to interpret policies? How are policies affected by different market models, e.g. property based, unlicensed, or brokered?
4. What language should be used to describe radio module capabilities? What radio interface should be presented to the application? How does one derive a common application interface from a specific module description?
5. How does one quantify mission requirements? How does one describe the capabilities of radio software modules? What techniques effectively translate from mission requirements to a radio configuration?
6. Given all the information available to a radio, how are operating settings derived?
7. How do multiple radios collaborate? What information is exchanged and what protocols are used? How are setting changes coordinated across the wireless networks?
8. How are each of these adaptive techniques secured against intruders?

These early research questions call for a active and rigorous research program. Investigators should be encouraged to develop and analyze the necessary measurement and adaptive algorithms. The most promising algorithms should be tested and evaluated in numerous testbeds and field-based experiments and integrated with the larger fixed infrastructure. The interactions between CR networks through the fixed infrastructure and with the fixed infrastructure need to be investigated.

Cognitive radios will operate in a dynamic radio environment. They will sense the environment, user demand, and radio performance and react to those measurements. Cognitive radios will exchange information with peers to use the radio spectrum efficiently, meet user demands, and achieve robust operation.

Impact of Cognitive Radio Research on GENI

Many of the research questions essential to building cognitive radio networks are, in some sense, extreme problems of wired networks. For example, both wired and wireless networks need to deal with links going up and down. However, in the wireless network, the frequency of link status changes is much higher than in today's wired network. So, wireless network architectures must pay closer attention to link status changes and react faster to these changes. This work would carry over into wired network.

Some characteristics of cognitive radio networks that are applicable to the larger, end-to-end network are:

1. *Operating environment sensing* – Cognitive radios measure and react to the environment they are operating in. This environment is multi-dimensional; including cooperative and non-cooperative emitters turning on and off, CRs adapting to their local changes, and traffic loads; and rapidly varying. CRs must rapidly adapt to this changing environment and communicate their changing operation settings to other wireless devices in the network. The mechanisms and techniques to sense, adapt, and communicate operation state are necessary in CR networks and applicable to networks in general.

2. *Robust communications services with unreliable links* – The radio links, by their very nature, have intermittent outages. A link outage may result from the temporary location of the receiver, transmitter, and other objects in the environment. CRs, by their very design, must deal with these very short term link outages, and do so through a variety of techniques. It is through this large set of techniques and mechanisms that wireless networks implement a robust and reliable communications service with unreliable links. The techniques and design patterns used in wireless architectures are applicable to the larger network architecture.

3. *Operational state languages* – CRs, as they adapt, must communicate their observations and operation state to other CRs in the network. A few "languages" will be needed to describe observations and operation state. This information is likely to be much richer than common link status information. For example, one radio might send a list of all emitters it has recently sensed to other CRs in the network. The entry for each emitter might include a frequency range, time, and spatial location, and signal format (e.g. spread spectrum or narrow-band FM). The language used to describe observations and operation state will be much richer than conventional node or link state information. The language(s) and protocols necessary for CR networks should influence general network architectures.

4. *Distributed Resource Management* – The radio spectrum is a distributed resource. Use of the spectrum in one location affects the availability of that spectrum in other network locations. Allocation of the radio spectrum resource must be carried out in a cooperative manner and balanced between (quick) local decisions and (optimal) global allocation. The algorithms developed to allocate the distributed radio spectrum and mobile network resources based on traffic loads and operating environment are applicable to the GENI infrastructure – and will require demanding new services within the GENI network.

These examples show how techniques and mechanisms necessary to CR networks will have an influence on the architecture, design, and implementations of networks in general.

2 Cognitive Radio Technology Roadmap

Cognitive Radios will advance as circuit technology and radio control and management systems advance. Today's primitive CRs are built from general purpose processors and common signal processors and field programmable gate arrays. In the future we anticipate rich developments in radio control and management, increased processing capacity at all levels, and improved distributed resource management. A technology roadmap showing anticipated developments in cognitive radio is given in Fig. 4.15 below.

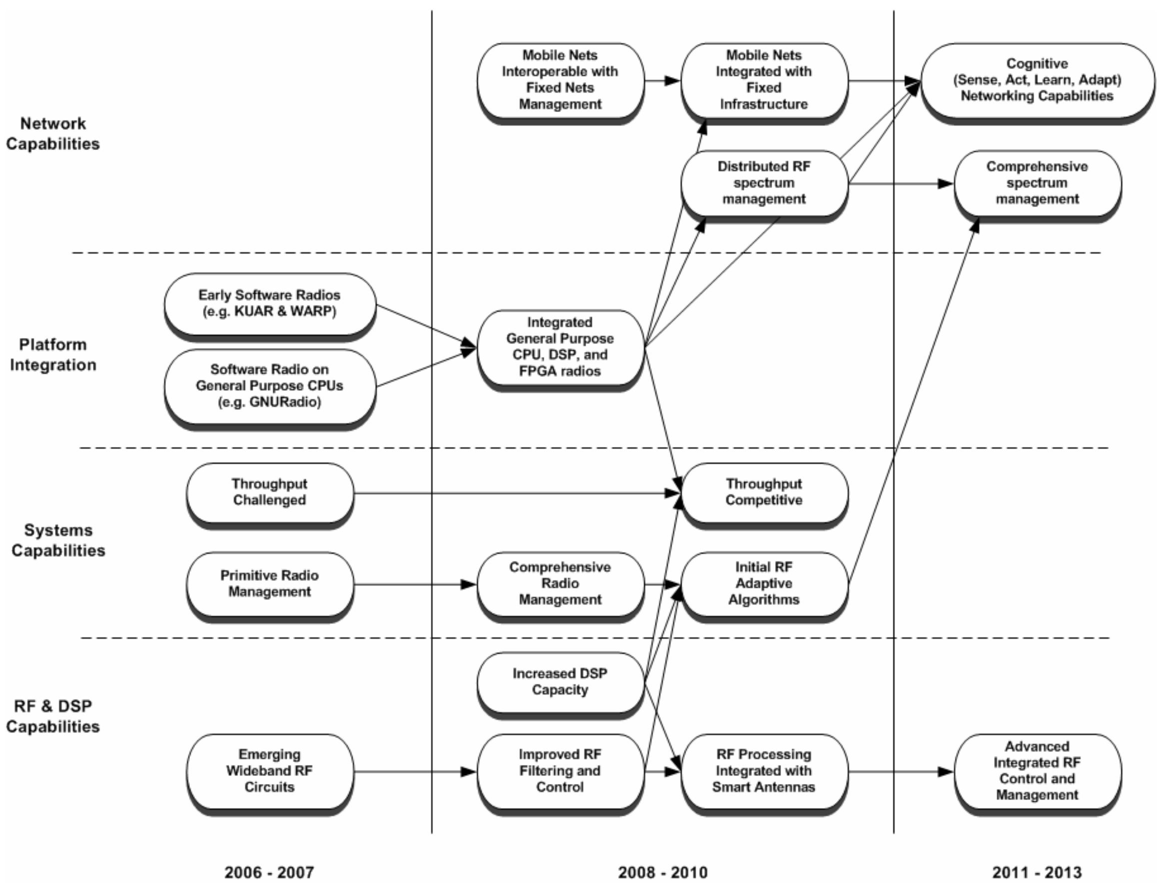

Figure 4.15: A roadmap showing anticipated developments in cognitive radios.

Several cognitive radio platforms are currently under development, some with NSF funding (for example, the GNU radio from University of Utah, the programmable radio kit from University of Kansas and the network-centric cognitive radio from Rutgers University, GA Tech and Lucent). These platforms are expected to mature and be fully tested for larger scale use by early 2007, and one or more of these designs will then be converted to medium-scale production necessary to implement this experimental system. A general-purpose cognitive radio platform includes reconfigurable hardware and/or DSP for radio modem signal processing, along with one or more embedded CPU's for packet-level protocols and radio adaptation/control. Such a platform is inherently suitable for slicing, virtualization and user programming as the same piece of hardware can be used to implement multiple radio technologies that co-exist in the network at the same time. Platform software for these emerging hardware implementations will be designed to take these requirements into consideration. Radio management will be limited and integration with the existing infrastructure will be at a basic level.

During the 2008 – 2010 timeframe we will see continued increase in signal processing capacity in general purpose processors, signal processors, and field programmable gate arrays. Improved radio management, adaptation, distributed resource allocation and interoperability with the fixed infrastructure will emerge. In 2011 – 2013 we expect to see comprehensive integration of mobile networks evolving into more adaptive, cognitive networks. Dynamic, high capacity mobile networks will be the primary means by which the majority of users access the network.

A critical technology bottleneck for cognitive radios is the performance of RF front-ends. Future multi-band RF front-ends will likely incorporate one or more flavors of MEMS components and integrate traditional front-end parts, such as analog RFIC devices with digital (DAC and ADC) / DSP / FPGA functionality. Most dedicated front-end RFICs (typically for cellular and 802.11 applications) use some flavor of ZIF (Zero IF), where the modulator / demodulator converts directly between baseband and the RF frequency range. Because of flexibility and volume reasons, research-oriented cognitive radio platforms may need to use a "hybrid ZIF design"; that is, capable off-the-shelf ZIF modulator / demodulator devices along with traditional mixer-based superheterodyne front-ends to convert from the ZIF device range to the target bands. A hybrid ZIF modem device, with embedded D-A and A-D converters, internal I & Q compensation, and digitally tunable baseband filters, essentially provides the capability to have "bits in / RF out" – and a few instantiations can provide coverage of many spectrum bands of interest.

3 Cognitive Radio Software

The cognitive radio outlined above requires an extensive software system to perform the basic functions identified above and emerging capabilities. In this section we provide a very preliminary description of a possible software organization for a cognitive radio. As cognitive radio technology develops, this organization is likely to be extensively modified.

Traditional telecommunication services attempt to divide functions into data, control, and management. It seems premature at this time to attempt to categorize cognitive radio functions into these categories. Hence, we will simply describe the major software functions.

The primary functions in this preliminary cognitive radio include:

1. *The radio hardware*. The radio hardware includes radio frequency circuitry and signal processing devices. We anticipate this hardware will be able to provide a description of its capabilities to enable self-configuration.
2. *Software modules*. Software modules represent code that has been loaded into field programmable gate arrays (FPGAs), digital signal processors (DSPs), or embedded general purpose processors. Since these are software modules, each defines its own interface to other software components. A common language to describe the interface would be useful.
3. *Middleware*. The middleware layer attempts to reduce the details of specific devices and software modules to common abstractions. For example, setting the transmit frequency of the radio frequency circuitry, or setting the encryption key of a software module. Creating a middleware system will require development of a common model for a wide range of hardware/software modules.
4. *Logical Radio Layer*: Depending on the radio configuration, the hardware and software can be programmed to act like multiple radio links are available. For example, the radio might support communications on several frequencies, time slots, or CDMA codes, each of which looks like an independent link. The logical radio layer implements this abstraction.
5. *Device Manager*: The device manager loads radio configurations into the hardware components and sets-up the logical radios.
6. *Configuration Manager*: The configuration manager determines which radio configurations are available on the physical radio for rapid loading into the hardware. It also interacts with module libraries (below) to determine which radio modules are needed to meet user requirements.
7. *Module Libraries*: The module libraries are collections of radio functions. For example, modulations (AM, FM, BPSK, QPSK, etc.), error control, encryption, and adaptive algorithms. The module libraries are build with a variety of tools (e.g. general purpose compilers, cross compilers, hardware design languages, and FPGA design tools). Coordinating the multiple sources that may go into building a specific module is a challenging task.
8. *Rules Engine and Policies*: Policies are used to limit the operation of the radio due to regulatory, geographical, or physical constraints. Policies should be usable independent of a particular radio. A “rules engine” is used to interpret policies. Device Managers, Logical Radios, Middleware, and hardware drivers might use the rules engine and loaded policies to determine allowed operation.
9. *Smart Controller*: A “smart controller” manages all of the radio resources outlined above. Currently, the Department of Defense’s Software Communications Architecture (SCA). In the SCA model, devices and software modules can register with the SCA system and locate other devices and modules within the radio.

Bringing software definition to radios opens a wide range of opportunities to enhance radio communications and wireless networking services. Along with that flexibility comes a challenge to reliability and robustly manage all the components that make up a cognitive radio. This preliminary organization and description above provides some insight into the research challenges of a cognitive radio. Several community and commercial efforts are developing cognitive radio software, such as GNU Radio and Vanu, Inc.

4 Experimental Infrastructure for Cognitive Radio Networks

This section provides additional detail on each of the cognitive radio network components outlined earlier in Sec. 3.5.1.

(1) A Flexible Cognitive Radio Network Test Environment

We expect future CRs to use a variety of hardware platforms and work in and over a very wide range of radio frequencies. During the design and development of these future radios, it may not be possible to operate the radios in the open environment. We propose one or more large, Cognitive Radio Test Environment (CRTE) to support research in emerging and future CR platforms. Characteristics of a CRTE include:

1. A large, shielded anechoic (RF) chamber – The anechoic chamber would be shielded (Faraday cage) from outside emitters (interference) and would allow innovative emissions inside and prevent emissions from leaving the chamber. The environment can be configured with movable absorbing and reflecting partitions and objects to create many different, but repeatable scenarios. The CRTE should be on the order of ~100’ x 100’ x 15’ at least and perhaps it should include two floors.
2. RF Emission Generators -- A number (~25) of signal generators would be available to place in the CRTE to generate pre-planned RF stimuli. These “RF scenarios” would be used to test and evaluate the responses of CRs to a changing RF environment. The generators would computer controlled, to easily define and repeat RF scenarios) and movable throughout the CRTE.
3. RF Measurement Instrumentation – The CRTE would include RF test equipment to measure “RF environment” for “reference” analysis. By recording the RF environment, which would include the RF Emission Generators and the CRs, a record of what each CR encountered would be available for off-line analysis. The RF record could be matched with measurements from each CR and the responses of each CR to the RF Scenario. The RF instruments would be computer controlled and movable throughout the CRTE.
4. Data Management Facility – The CRTE should be computer controlled. This means it should be easy to design the layout of a specific experiment, define the RF Scenario for the experiment, define the measurements to be collected, and analyze the results. Experiment results should be available to the research team and the research community. This facility would be a set of computers and databases.

(2) Open Air Cognitive Radio Network Test Environment

A second facility is a set of (~100) SDRs arranged around a city or town that would be available for testing CR algorithms "in the real world" – see Fig 10. The SDRs would be linked by a wired infrastructure and necessarily operate on unlicensed frequencies or "experimental frequencies". The Open Air CR Testbed (OACRTE) provides a facility to test and evaluate CR concepts in a more realistic environment. It also provides a mechanism to evaluate inter-operation with other technologies, such as vehicles, sensor networks, future cellular infrastructure, and the evolution of IEEE 802 technologies.

OACRTs would be deployed during the later stages of the project in one or more locations. Locations would reflect typical wireless service environments, e.g. a suburban area and an urban area. A cognitive radio deployment in a suburban/medium-density coverage area of approximately 50 sq-km will support demonstration and evaluation of this technology as an alternative to available cellular and hybrid cellular/WLAN solutions. Facilities such as the Table Mountain National Radio Quiet Zone might also be appropriate for CR experimentation.

Implementation of an open air system will involves construction of a distributed spectrum measurement infrastructure along with centralized spectrum coordination resources (such as spectrum broker, spectrum server). A new wideband experimental spectrum allocation will also be required to support this trial network. A total of ~50 cognitive radio routers will be deployed over the coverage area. The deployment will also include ~500 cognitive radio terminals (associated with end-user applications). End-users will be able to program their own layer 1 (radio physical layer), layer 2 (data link and medium access control) and higher protocols on these devices through the experimental software interface provided by GENI.

Researchers would be able to reserve time on OACRTE to run experiments and collect measurements.

(3) Software Defined Radio Experimental Platforms

Software Defined Radio Kits (SDRKits) are a mechanism for easy, widespread deployment of experimental radio systems. SDRKits enable research groups without the capability or time to develop their own platforms to participate in developing, testing and evaluating Cognitive Radio concepts. A set of a few common radio platforms will enable significant sharing of radio designs among research groups. Further, a few common SDRKits will make experiments in the CRTE much easier and support deploying significant numbers of radios in urban-scale experiments. Programmability ensures that the deployed infrastructure can evolve over time.

SDRKits feature programmability at many layers: protocols for network management, routing, and congestion control can be programmed in conventional languages. Functions that require high-performance implementation are programmed in a hardware description language (e.g. VHDL or Verilog) for execution on FPGAs. SDRKits would also have standard interfaces for connection to legacy protocols such as Ethernet and WiFi, and custom interfaces for high-speed interconnection with custom wireline GENI infrastructure.

Programmability at the lowest layers, i.e. signal processing logic, is critical to understanding and obtaining real-world experience with real-world radios and applications. The performance enabled by directly programming FPGAs provides the high data rates that wireless users expect. Likewise, FPGAs will enable rapid reconfiguration of a large-scale, already-deployed testbed. As wireless technology changes over the course of GENI, these changes can be rapidly incorporated into field trials.

SDRKits can be used for large-scale urban deployment and small-scale laboratory experiments. In particular, many researchers will obtain a moderate number of kits, order 10, for local in-laboratory or small-scale outdoor prototyping and demonstrating proof-of-concept. After initial local testing, the same radio configuration can be used on larger-scale GENI facilities with significant interconnection to the GENI infrastructure. The large-scale facility will support extensive end-to-end experiments.

SDRKits for GENI should meet three key requirements.

1. SDRKits must directly support development of innovative technologies in multiple interacting areas that include physical layer algorithms, data link protocols, media access, and radio system architectures, among others. The hardware should be scalable to support algorithms of variable complexity and should be extensible to support a diverse array of I/O devices. Further, the platform should have an upgrade path to ensure the continued development of technology.
2. SDRKits must support experiments ranging from individual laboratories to full end-to-end systems. Further, kits must facilitate fully operational open-access networks at every level of granularity.
3. SDRKits should have industry input and support to guarantee industrial-strength manufacturing, software tools and product support. Simultaneously, it must be inexpensive and widely available to academic and research institutions for research, education and deployment.

(4) Simulation and Emulation Tools

Cognitive Radios operate in a complex physical environment. However, evaluating radio performance in a physical environment is expensive, difficult, and researchers generally must be present at the site. Simulation and emulation tools enable a larger group of research teams to participate in the development of CRs. We are proposing a facility that will support simulation and emulation of CR networks of 10’s to 100’s of nodes. One to three facilities around the country will support multiple simultaneous users, allow interaction and interoperability with the larger GENI infrastructure, and provid professionally supported tools.

The Cognitive Radio Simulation and Emulations Test Environment (CRSETE) will be integrated with the CRTE and the OACRTE. Measurements from the physical test environments will be archived and made available to refine simulation and emulation models. Common simulation scenarios describing radios, environments, mobility, infrastructure, and applications will be developed to enable researchers to compare and evaluate their designs and implementations.

Test Environment Summary

We have described four GENI capabilities that support research and development of advanced and innovative wireless networks. The CRTE supports those research teams developing adaptive radio algorithms and distributed radio spectrum resource management methods. The OACRTE supports testing CR implementations in a real world setting. SDRKits, the OACRTE, and the CRSETE support those research teams that are interested in developing cognitive radio technology, but do not have the facilities, capabilities, or time to develop their own platforms. This range of GENI capabilities offers a significant opportunity for inter-disciplinary research and advancing the state-of-the-art. The proposed facilities are also an excellent place for graduate and undergraduate students to gain experience in advanced wireless networking.