COGNITIVE RADIO SYSTEMS

Dr. Hossein Eslambolchi

March 1, 2012

INTRODUCTION:

This report reviews the major trends in a variety of key technology areas of potential importance to Ericsson. Included are discussions of technology markets just emerging, such as ultra-wideband radio, as well as establishing technologies with important technical innovations, such as Web Services and Grid computing.

The chapters are intended to provide a brief overview of the subject technologies.

This Report is intended for Ericsson Inc. to drive new innovation to build an eco system for convergence of IP and Mobility.

Contents

Abstract:

This white paper discusses the notion of “cognitive” radio, the combination of a radio transceiver with computerized intelligence to automate coordination of devices, networks, and services for improved functionality, interoperation, and spectrum utilization. The term originates from parallels with biological systems that interact with their environment using a regime of goals, sensory inputs, and reactive behaviors to sustain themselves.

Cognitive radio, formerly the interest of a few, has now entered the “group think” environment, where “cognition” means many things to many people. Outgrowths of the concept have been advanced to address a variety of next-generation wireless challenges. The paper examines the evolutionary possibilities, as well as tracing its origins and current forms. The author also attempts to forecast how FCC cognitive radio rulemaking may proceed in connection with its initiatives to achieve higher spectrum productivity.

What Is Cognitive Radio?

Like many modern concepts, "Cognitive Radio[1]" is not entirely new, but rather an amalgam of manual radio resource management techniques invented and proven earlier for enhancing the performance of wireless communication systems, now automated by a computer. Simply put, cognitive radio is the combination of a radio transceiver with computerized intelligence to automate coordination of devices, networks, and services for improved functionality, interoperation, and spectrum utilization.

Recently a focal point has become coexistence of radios in common spectrum. It is termed "cognitive" because the actions it is envisioned to take are analogous to the way a biological entity reacts to its environment based on a regime of goals, sensory inputs, and behaviors. Cognitive radio has become a topic of interest lately because of two important changes in the radio application space:

1. The growth of personal portable communications services
2. The need to provide enough suitable spectrum for such services

Personal communications has changed the face of two-way radio from a microwave relay, push-to-talk, and mobile telephone world to mainstream, anytime, anywhere wireless for the masses. Cellular telephone started this revolution in the 1980s, but today's broadband multimedia wireless networks have carried the banner forward.

Even at its inception, the very nature of cellular radio dictated some level of cognition: The systems depended upon devices to detect an available wireless network, interact with it to allow a connection to be set up, and to maintain the connection as the user moves about. It is fortunate indeed that cellular and microcomputer technology arose contemporaneously, as much of what cellular telephony has become would have been impossible otherwise. In today's advanced radio networks, even more is expected: wireless devices may actually be parts of the network as well as clients (as with forwarding mesh topologies), they may have to accommodate a variety of media, rates, and

QoS needs, and they may have to operate in spectrum that is shared with other services while minimizing interference. They may even have to recognize exactly where they are and to notify other devices or networks.

The Origins of Cognitive Radio

Having begun almost 100 years after the development of wired technologies, radio has traversed several stages of maturation, the latest of which is exploitation of the signal processing techniques that had proved so successful in making broadband transmission possible using wires. However adoption of these techniques has also forced radio engineers to reconcile with Shannon: Information theory increases consciousness of the bandwidth resource, link attenuation, power, and noise floor as fundamental limits of reach and rate. This entry into the "Shannon Zone" has also sparked the realization that spectrum is finite and not all of it is necessarily appropriate for every application.

The realization that spectrum is finite is easily visualized if one thinks of a wireless "cell" as a coaxial cable of the same diameter. Individuals within the cell/cable must share the RF medium efficiently with others while minimizing interference. With wires, one can always add more cables to fixed-use locations because the spectrum is contained and can be reused almost infinitely; with wireless the spectrum is much less contained and can be reused only to the extent smaller cells can be used. Even with small cells, the spectrum resource may have to be divided among a significant number of users within the coverage area. With appetite for wireless broadband connections increasing, cells must inevitably become smaller with better resource reuse and higher channel spectral efficiency absent vast amounts of new spectrum.

In the end, one might envision a "Shannon Communication Volume" (SCV) expressing the maximum number of users at a particular rate that can be sustained over a geographic area in much the same way as the Shannon limit is now routinely used to bound the "reach" and "rate" of transmission over wires.

The Communications Act of 1934 established the FCC as the manager of spectrum for citizens of the United States and a regime of licenses, regulations, and process thereafter ensued to ensure that spectrum use was organized, productive, and efficient. In a sense, the FCC's goal has been to maximize the SCV of the United States by building rules for various radio services one-at-a-time into a large body of regulations.

Smart Radios:

Over the past quarter-century the marriage of wireless and inexpensive computing has allowed "smart" radios to operate within the "one-at-a-time" FCC rules to make spectrum work harder. Since smart radios can do more than simply receive and transmit radio signals---they can also measure them and respond to propagation conditions or availability of known services---they are at the very heart of modern radio resource reuse concepts. Many of these capabilities were pioneered by 1st generation analog cellular systems, wherein phones would search for an "overhead" channel, adjust their power level according to information provided by the base station, and then register" with the system to use it. Because intelligence was costly then, much of it was centralized at the MSC. As microcomputer technology became increasingly capable and economical, more and more of the intelligence migrated into the terminal. Soon, transmit power was adjusted dynamically according to received signal strength. Still later, Mobile Assisted Channel power was adjusted dynamically according to received signal strength/ Still later, Mobile Assisted Channel power Assignment (MACA) and Mobile Assisted Hand-Off (MAHO) used signal strength information collected by the mobile radio itself to determine the best serving channel or handoff target. Likewise, early digital cordless telephones could detect interference at the handset or base and retune or change hopping sequences to improve operation.

In spite of the ability to react to signal strength, even 2G digital cellular systems were designed for only one connection type: voice telephony. Accordingly, they could deliver only one "lowest common denominator" service level that,

that which could be sustained at the edges of a cell. This constraint is the basis of today's 2.5G cellular service, even though "data" packets may now be transmitted on a "voice" channel.

"3G" cellular and 802.16/WiMAX smart radios have moved to the next step: Link adaptation. Increasing experience with digital transmission and use of more sophisticated modulation, coding, and antenna techniques have taught that if users are close to the base station they can probably sustain higher rates than those at the cell edge. Using improved link measurements, the radio can adjust or "gearshift" its rate, coding, and spatial signal combining to optimize performance for the particular propagation environment in which it finds itself. This improved "cognition" is an integral part of 3rd generation systems, and the adaptation has lately extended to advanced antenna technologies such as MIMO and beam-forming, providing a new kind of intelligence---"smart antennas"--- for smart radios.

SMART MAC RADIO:

In the 1980s, a new standards creation strategy was popularized: the ISO OSI[4] transmission model. This model sought to break a communications path into individually specified functional "layers", each with identified inputs and outputs to other layers or to the transmission medium itself.

Prior to the ISO model, most radios were designed holistically with all functions interoperating broadly and capable of sharing information freely. Most mainstream wireless communication systems (*e.g.* cellular and cordless) also used frequency-division multiple access and frequency duplexing to share the radio resource in the spectral domain. The "paired" nature of spectrum allocated for these services made use of arbitrary spectrum

The ISO Model Compartmentalized" intelligence, an specific of the format impractical. of all messages that traversed layers, in effect producing "smart layers". It simplified standards setting and device interoperability, but limited the ability to communicate the products of each layer's intelligence to other layers. With physical layer (PHY) transmission "machinery" in place and control messages specified, most intelligence eventually rose to the MAC layer. The flexibility of packet transmission combined with the ISO model, the "smart MAC", and the reduction-to-practice of time-division duplexing opened the field for communication services in non-paired spectrum. Now it was practical to move services to new bands on an opportunistic basis.

802.11 was the first standardized packet radio framework to capitalize on the opportunity and simultaneously address the problem of allowing users to share a common channel asynchronously, and its appearance was just in time to "catch the wave" of Internet popularity. It borrowed a technique for mediating transmissions adopted from Ethernet, which had previously capitalized on the shared-medium packet upsurge: Carrier Sense Multiple Access – Collision Detection (CSMA-CD). However since 802.11 targeted not a wired application, but rather wireless, collisions could not always be detected as with Ethernet. Accordingly, the protocol was augmented to implement Collision Avoidance (CSMA-CA) instead. CSMA-CA uses Request-to-Send/Clear-to-Send messages to effect "reservation" of the medium that discourages access by other stations until the sending station has completed its packet transmission.

CSMA itself arose from earlier ALOHA[6] protocols developed to allow packet transmissions within a network of radio stations with completely decentralized control (very unlike cellular systems). With the simplest ALOHA protocol a station completely transmits its entire packet when it arrives. If a collision happens with another station that has also transmitted (signaled by not hearing a repeat of the message from the intended receiver), the station retransmits the entire collision-damaged frame after a random delay. ALOHA could be termed "strongly challenged" on the cognition scale, and its efficiency was low. Various schemes (including "slotted" ALOHA) subsequently improved its performance somewhat at the cost of synchronization. However with many active users, ALOHA networks experience "congestion collapse" where throughput (and system efficiency) may be reduced to near zero.

Simply put, CSMA extended ALOHA by implementing true spectral etiquette for packets including "congestion awareness": if a device has a packet to send, it listens to see if the channel is in use. If not, it sends the packet. If it fails to receive an acknowledgement, it tries again after a congestion-modulated random delay to send a duplicate of the damaged packet.

In adopting CSMA for Media Access Control, 802.11 developed not only a reasonably efficient means of utilizing an isolated cell's capacity, but also created a simplistic method of sharing the radio resource between groups of 'embedded" cells using the same frequency whose coverage areas might overlap. The multi-cell sharing property was enabled by the so-called "Clear Channel Assessment[7]" function that determines whether the MAC can transmit with reasonable assurance that it will not harm transmissions already in progress. CSMA enabled easy assembly of networks that could cover larger areas than a single WLAN cell, and paved the way for "do-it-yourself" networking. Use of CSMA for cognitive spectrum use is now lumped under the heading of "contention based sharing techniques".

Private do-it-yourself wireless LANs capitalized on the contention-based protocol advantage, while contributing another biological-like behavior to the cognitive radio vision: the "organic" growth model where a network of small cells expands coverage and capacity by "birth" of new "smart" base stations that "bond" to the old ones fueled by the presence of nutrients (*e.g.* users, utility, investment, *etc.*). The potential of such "self-organizing" networks is a fundamental driver of 4th generation broadband wireless thinking, since the sheer number of cells required to provide true Ethernet rates and contiguous coverage to large areas mandates a distributed-intelligence control architecture for scalability and cost reasons. But perhaps equally importantly, the small-cell precept also makes very dense radio resource reuse practical, particularly in spectrum that is either unlicensed or licensed in less-rigid ways.

Although WiFi networks covering large areas (even entire cities) have lately become fashionable, there is an important limitation that must be overcome before such networks can truly be considered broadband multimedia systems. This is because 802.11's implementation of CSMA works well for bursty, non-time-sensitive data packets, but not so well for streaming real-time services such as VoIP. The reason is that the latency CSMA systems develop as they become more heavily loaded can cause packets to be received too late for decoders to process them as part of the stream reconstruction process, causing loss of quality. 802.11 addressed this shortcoming by creating the 802.11e standard, which establishes two new operating modes to support QoS-dependent applications. The most simplistic conceptually is EDCA[8], which allows QoS clients to use shorter back-off times. The approach provides “prioritized” QoS in lightly-loaded systems, but cannot guarantee QoS when systems become more heavily loaded. The second mode, HCCA[9], is capable of providing scheduled radio resource management by suspending CSMA operation in part of the 802.11 super-frame, while instituting a QoS specification, scheduling, and polling process that provides reserved “Transmit Opportunities[10]” for clients requesting “parameterized”, carrier-grade, QoS. Since operation in the contention-controlled portion of the frame provides a “protected” radio resource, a “smart scheduler” can materially improve system performance, not only for clients, but for better reuse in embedded cell environments.

The Smart-Scheduling Radios:

The simplest definition of a wireless MAC is a protocol that tells radios when they can transmit. For networks where no centralized admission or management of clients is used---so-called peer-to-peer operation---clients all transmit packets according to CSMA-like rules. However for “infrastructure-mode” networks in which APs11 connect to a backbone, MACs operate differently according to whether the device is a client or an AP. As has been discussed above, 802.11e HCCA depends on strong scheduling coordination to minimize collisions, allocate TXOP lengths sufficient to maintain QoS for individual clients, and maintain high network resource efficiency. One important way this can be of value is associated with the so-called gear-shifting transmission behaviors already touched upon in previous sections.

Gear-shifting is a term used to describe adaptive behavior that tailors transmission rate to link quality. In a smart radio, the need to shift “gears” is usually triggered by examination of error rate and signal strength as a session proceeds. If error rate increases or signal strength decreases, the MAC responds by transmitting data more dependably, but also more slowly. If conditions improve, the MAC “shifts” to a higher transmission rate. Gear-shifting can be viewed as “local” adaptation, because it improves the performance of a particular client being served by the system. The consequence of the relationship between energy-per-bit and throughput with fixed transmitter power and a given link attenuation means that a "downshifted" client's application gets fewer packets over a period of time. The optimization allows that client to receive some packets as opposed to getting few or none at all. With bursty, best-effort data applications, the slower behavior causes applications to appear less responsive (e.g. Web pages will "paint" slower), as those who routinely use dial-up modems know. Absent gear-shifting however, long outages due to errors could cause TCP/IP protocol timeouts, possibly closing sessions.

But streaming real time traffic is a different situation. These streams connect applications that must produce synchronized time-critical media outputs at each end. In QoS enabled systems, clients are admitted according to a TSPEC12 that attempts to furnish a minimum rate with a maximum latency and jitter. Since longer bursts are required to maintain the same throughput following a "downshift", more time must be allocated for these clients to allow them to stay synchronized in real time.

In a distributed coordination system (such as EDCA) we may illustrate the effect of local optimization of one link upon other links, showing that gearshifting can become a global resource allocation problem: A large number of QoS-bound streaming clients is being supported under stable conditions. One client moves away from the serving AP toward the edge of the coverage area. The signal strength at the client and AP drop, and the error rate increases. Since at the edge of coverage, radio error rates rise strongly as the signal at the receiver becomes noisier, the AP attempts to compensate. It gearshifts to accommodate the poorer conditions, transmitting longer bursts to the more distant client in an attempt to maintain TSPEC throughput. The throughput for other clients must now be accommodated in less time, raising apparent system loading. Contention builds due to the higher loading rate, causing packet queues to lengthen and latency to rise. If more clients force downshifts or more clients enter the network, the AP may then not be able to support new streams, accept incoming handoffs, or continue to honor admitted QoS specifications.

With a point coordination system using a strong scheduler such as that associated with HCCA operation, clients are admitted only to a certain resource-use level. Time is held in reserve to allow some QoS clients to downshift their sessions as circumstances require, while still allowing other clients to be admitted either as new sessions or due to handoffs from other cells. If the aggregate throughput in process reaches a pre-set policy limit, it is possible for an entering client to experience a "spectrum busy" indication. The client can respond by waiting and trying again, or generating a TSPEC with lower rate or higher latency.

In such an HCCA system, the smart scheduler becomes the heart of a smart MAC radio: it dynamically adjusts session TXOPs, admissions and the reserved resources to globally optimize network resource efficiency as well as locally optimizes individual client performance---up to a point. Since wireless is by definition un-tethered, it is possible that a client could request downshift that might require many times more TXOP length than could reasonably be accommodated at the current network loading.

If the amount of best-effort traffic is insufficient to support "capacity borrowing" without violating policy minimums for best-effort service, the downshift may have to be disallowed. Under such conditions, the scheduler may trigger a notification that the end of range has been reached, and that the stream cannot be supported if the link were to deteriorate further. Seen in this way, the smart MAC becomes the globally-aware arbiter of when downshifts are needed, if they can be accommodated, and the guarantor of QoS when the downshift is complete.

By extending these concepts, a scheduler can also be used to coordinate APs as well as clients because it can allow several device classes to operate in a super-frame simultaneously. This property can be used to substitute for the ability of CSMA to support resource sharing in embedded cell layouts where two APs may be operating on the same frequency, while simultaneously maintaining strong QoS. In such a case, the schedulers of two APs that detect they would interfere with each other can mutually-negotiate exclusive time allocations in the super-frame. They can then each manage an individual segment of the contention-free period using their smart schedulers.

This time-sharing allows smart MACs/schedulers to create an optimizing environment with homogeneous radios in which APs (as well as clients) can share the radio resource in time, frequency, and space to form a reuse pattern dynamically. The concept, of using schedulers to perform this function was developed at Lucent Bell Labs – Research, and is called “CelluLAN”, since it fuses cellular frequency-only and wireless LAN time-based reuse strategies. We shall further explore these concepts when we discuss “Spectrum Aware Radios” in a later section.

Smarter Radios: Digital Radio Processing

A natural extension of smart radios and smart MACs is the Digital Radio Processing (DRP) concept. DRP extends the perceptual and interactive capability of the radio beyond that of a single air interface. Before DRP, the “front end” of a wireless device was designed with components specifically tailored for a particular radio service, and hence a particular frequency band, channel bandwidth, modulation, error correction code etc. Consequently creating a single device capable of accessing several systems became a difficult integration problem: Essentially, separate radios had to be “shoehorned” into a housing with the computer selecting which one is to be used at any point in time.

DRP uses a different approach. It uses very high speed computing to operate on a large segment of spectrum, making few assumptions about the signals that may inhabit it. To receive signals, a wideband RF “front end” captures the segment, amplifies it, and translates (heterodynes) it to a lower “intermediate” frequency. A high-speed, high-resolution analog-to-digital converter is then used to sample the waveform represented by all of the signals in the segment and convert it into digital form. The digitized samples are then numerically processed to mathematically realize the channel tuning, filtering, and information-extraction functions normally accomplished by analog means. The reverse of the process is used for transmission: the desired waveform is synthesized digitally, converted to analog samples, filtered, translated to the final output frequency and amplified before being fed to the antenna.

The numerical calculations required to implement DRP can be accomplished by using one of two computing means: an “engine” or a “general-purpose programmable DSP”. The “engine” approach, sometimes called a “Programmable Hardware Radio” utilizes a group of specially-designed signal processing subsystems on a VLSI chip to emulate a local oscillator, mixer, filter, and demodulator (or the corresponding operations required for a transmitter). Each subsystem can be programmed parametrically to adjust its behavior as desired. As the mathematics needed to realize radio functions does not change when parameters are reprogrammed, the engines to do the job can be realized using an ASIC13 design very effectively. Such an approach has the advantage of power efficiency, because each subsystem is designed for maximized throughput per watt of dissipation. Because new communications systems use sophisticated signal processing already, such as for OFDM transmission and detection, many “baseband” ICs already use similar techniques, albeit at lower sampling rates. Examples are the chipsets that comply with 802.11a, b, and g standards, which all use the same channel bandwidth, but operate with different RF frequencies and modulation techniques. These applications provide DRP-like processing power, but are not field-configurable as a "real" DRP radio would be.

The second DRP approach is less power efficient, but does not require special purpose VLSI. In this variation, sometimes called "Software Defined Radio" or SDR, a DSP executes a program to perform the mathematical computations. Although this "all software" approach is compelling, the DSPs must be very fast to keep up with the incoming samples, as they must complete many more internal machine cycles to process them. Because of the inefficiencies, the approach tends to be less attractive for battery-operated equipment. With seemingly relentless computing progress driven by Moore's Law, we may expect that SDR will eventually replace engines.

However, battery limitations will remain an important consideration for portable DRP applications for the foreseeable future.

Regardless of whether ASIC or SDR processing is used, DRP can materially improve the performance of a radio. The performance of a
digitally processed channel filter is shown contrasted with a typical "hardware" filter. Digital filters provide more-abrupt out-of-band attenuation, coupled with highly controllable, repeatable, in-band characteristics. This makes a radio much more capable of rejecting adjacent channel interference while separating the desired signal with minimum distortion, a particularly important consideration for modern dense-constellation digital waveforms. Digital filters are equally valuable in a transmitter for limiting modulated signal bandwidth and preventing "splatter" into adjacent channels.

The SDR Concept, As many modern wireless techniques, had its roots in military communications, where several different radio air interfaces might be used in used in favor force war fighter scenario, and where interoperability between battle forces is a valuable asset. SDR is also useful for monitoring clandestine radio transmissions because fewer assumptions about the nature of particular signals need be made. Perhaps the ultimate SDR application is the SETI, or Search for Extra-Terrestrial Life, project where the assumptions regarding a candidate "intelligent" signal may be based only on small differences between it and surrounding noise.

Ericsson Labs – Research has been a leader in researching use of DRP for commercial cellular applications, and developed radio architectures suitable for reception and transmission of a variety of cellular standards using the same device.

DRP base stations such as these were used to implement a self-organizing "underlay" system for to provide improved cellular coverage indoors. We shall discuss such concepts in more detail in a following section.

Multi-Service Radios:

The combination of several integrated radios or a single DRP radio under microcomputer control is an important new type of cognitive radio: Multi-service. In current cellular systems, microcomputer intelligence is used to establish communications using a single radio and Common Air Interface (CAI), while establishing a single application/user interface. The computer in a multi-service radio additionally detects and chooses which services are available in a particular location, provides a multi-service user-interface, and selects the network based on user or application-directed preference.

This is particularly important because a variety of wireless connectivity options are arising, augmenting cellular as the sole choice. Now, depending on the location of the user, two or more connection alternatives might be available. An example of this option is the appearance of handsets that are both cellular and WiFi compatible. With early products, connection to WiFi uses a separate radio in addition to the one used for cellular connections, leveraging the economy and low dissipation of 802.11 chipsets. However SDR capabilities using multiple RF front ends could enable 2.5G, 3G, WiMAX, WiFi, and 4G wireless connections (LTE) using the same processing hardware.

The ability to access several networks with the same device is harder than it looks, though: Multi-service radios must provide a user interface "cocoon" supporting intelligent connection to one of several access alternatives based on system availability and decisions derived from a rule-base that is part-programmed and part-learned according to user behaviors. A multi-service, multi-band cognitive radio may also have to orchestrate transfers between very different systems and services as a user moves, but without making "mistakes" that cause sessions to drop or produce handoff dither.

Re-Thinking Cognition and Spectrum:

Thus far, this paper examined cognitive radios that use computerized intelligence to improve individual system and service performance. They use "homogeneous[14]" radios with coupled FCC spectrum allocations, such as cellular telephone and wireless LANs. Recent spectrum-use studies by the National Institute of Standards and Technology (NIST) and others have revealed that the FCC allocations created since 1934 harbor many areas of frequency and geography where radio spectrum is underutilized. This is partly because individual allocations would have to have been too detailed to administer and too hard to police if every possible location, application, and interference threat were to be specified. Such studies also revealed that the purposes for which some allocations were made have become obsolete or have been changed by new technologies with consequent retirement of many of the original radio devices. This so-called "white space" has become a target of increased interest as the popularity of personal radio communication is overtaking broadcast and other wireless applications, and is driving the search for new spectrum to "wireless" America. As with all finite resources, which are usually exploited greedily at first, it is now necessary to revisit spectrum use to see if additional value can be extracted.

Within the last few years, attention has turned toward using cognitive radio as a means of augmenting the FCC regulatory role with "micro-management" of spectrum at the point of use. Clearly, a fully "sensory" radio might be able to understand enough about its local environment to identify other signals, gauge interference potential, and either "set up shop" or "move on to greener pastures". The enthusiasm for such capabilities was heightened by the increasing use of unlicensed frequencies to provide networked services.

Early FCC "Part 15" regulations for unlicensed emissions arose from the interference potential of reception by-products such as local oscillator radiation in receivers. However, responding to the experimental nature of radio, the FCC also allowed engineers to craft devices that could serve useful purposes without significantly interfering with licensed services as long as their "reach" remained quite limited. Examples included early wireless microphones, portable AM and FM "broadcaster" toys, *etc.* Eventually, some of these devices became capable of establishing two-way links, wherein they entered the "personal communication" application space.

Citizen's Band radio, which operated in a small slice of spectrum near 27 MHz, was established as a licensed service allowing 5 watts of RF power. It revealed steadily-building pressure for personal portable communication, but in the process became a de-facto unlicensed band famous for misuse. Pandora's Box had opened only a

crack, but far enough to expose the chaos that might result from uncontrolled use of high-power, long-reach unlicensed devices. The popularity of analog cordless telephones fully opened the Box.

The initial popularity of cordless phones not only stretched the limits of unlicensed operation and spectrum availability, it drove the FCC to open new initiatives to promote use of newer radio technologies (some pioneered by military applications) that could utilize spectrum more efficiently and be more tolerant of interference. To encourage these directions, it released spectrum in the so-called Industrial, Scientific, and Medical (ISM) microwave bands in the mid-1980s allowing transmitter power outputs much larger than had been possible previously in the unlicensed environment (except for Citizen's Band) The higher power tantalized equipment makers with the opportunity to achieve longer range if they used a then-new concept, "Spread-Spectrum". Spread-spectrum also promised improved spectral sharing with embedded uses (*e.g.* industrial RF processes, microwave ovens, and diathermy). Although not strictly a part of the new ISM regulations, the FCC encouraged device makers to observe "spectral etiquette", the idea of observing the frequency to determine if another device was already communicating before transmitting. Since most communication applications at the time were circuit-switched in character, devices would usually listen during idle periods between calls for channel "occupancy" and use the most inactive one should call setup be required. Some more sophisticated devices could also detect interference via error monitoring during a call, in which case they would attempt to switch channels or spreading sequences to improve quality.

The ISM regulations were an early attempt to allow various applications to cooperate in spectrum by equalizing interference "effects", particularly for the two very-different forms of spread-spectrum allowed, Direct-Sequence (DS), and Frequency-Hopping (FH). It was thought that by balancing DS and FH parameters, the two techniques would offer mutually equal communication disruption due to interference, while also controlling interference to lower-power, non-spread, systems.

The experiment was a wild success, bringing digital cordless phones and many other new devices to market. However, it also taught some hard-learned lessons: (1) Interference potential and noise are very hard to quantify reliably and even harder to respond to dynamically, and (2) peer-to-peer communication devices that succeed in the marketplace eventually spawn a desire among users to be networked (a variation on Metcalf's Law[15]). One might say the second realization predicted a need for "smart" radios to be "sociable" as well, prompting a number of "networked" cordless phone outgrowths such as CT2 and PACs.

The learnings from unlicensed applications and the desire for personal communication prompted FCC reconsideration of how spectrum could be better regulated, taking into account "application-based" (service-driven) rulemaking to augment the previous "device-based" mentality. This thinking was reinforced by growth of 802.11-like WLAN networks and by grassroots networking experiments such as SF-LAN. Cognitive radios that could be "spectrum-aware" appeared to open the way.

The Spectrum Aware Radio:

What makes a smart radio into a spectrum-aware radio? Principally, this is an acknowledgement that the radio will be used in a service environment where other radios may also be operating in the same spectrum, and that it must cooperate with them to use the spectrum productively. SA radios implicitly understand that some signals might be "foreign" (*e.g.* not using the transmission framework chosen by the network they wish to establish). Since foreign networks or devices could use different channel bandwidths, modulation types, MAC protocols, *etc.*, the SA radio must be able to act as a spectrum analyzer, spectral and time signature identifier, and interference-avoidance manager.

An interesting early example of a spectrum aware radio was the CDPD[17] system which could be called the first spectrum-aware radio implementation. Employed by a licensed cellular Provider like Sprint, this system operated in the same spectrum as 1st generation analog cellular telephones. It used a special SDR "sniffer" receiver to detect unused voice channels in real time without direct connection to the cellular base station equipment. A companion packet transceiver at the base station, also separate from the cellular system, exchanged packets with user data modems over a specific unused channel designated by a CDPD "beacon" until the channel was assigned to a cellular call. The CDPD transceiver would then move to another "sniffed" unused channel, mark it with a beacon, and resume operation.

One of the first applications of SA technology in public bands was to meet the requirement that 802.11a radios be able to scan, identify, and avoid certain radars used by the military. This proved to be more difficult than it appeared at first, since it was later learned that some characteristics of the radar's transmissions might change to maintain stealth. It also revealed that modern spectrum-aware, high-performance radios might have outgrown the ISO layered standardization model. This realization has driven disassembly of the layered approach in favor of an integrated PHY/ MAC/Network architecture, that can support higher rates, more sophisticated antenna processing, smoother gear-shifting, and enhanced spectrum awareness.

A major new opportunity for SA radio is operation in bands where older fixed radio services operate mostly continuously. This is because a significant amount of "white space" has been shown to exist in bands inhabited by legacy radio services with easy-to-identify emission "signatures". Such services were created under FCC rules that created licenses with inherent spectral or geographic "buffer zones" to guard against interference or a limited number of services to the exclusion of others. Frequently the licenses covered relatively large coverage areas by today's standards.

The availability of this newfound "white space" is particularly important since not all radio spectrum is suitable for mobile or portable applications. Microwave frequencies beyond ~10 GHz are so significantly attenuated by trees and objects that they are usually suitable only for line-of-sight operation. Spectrum below 50 MHz is fragmented, unsuitable for large channel bandwidths, and may exhibit sporadic long distance propagation behavior that complicates spatial reuse.

Consequently spectrum from about 50 MHz to 10 GHz is "beachfront property", and particularly valuable for so-called "underlay" operations, where communication applications could share spectrum in areas where the large-coverage area "buffer zones" or frequency-reuse layouts permit. For example, geographic distribution of larger-cell or point-to-point systems could allow small-cell under-lays to use frequencies deployed elsewhere. This spectral "leasing" concept is one thrust of recent FCC consideration on new rules using cognitive radio.

As we shall see, the "species" of SA cognitive radios used as a supplement to FCC regulations relate to the difficulty of the decisions they are expected to make, how much information is available, whether they are located in a spectral environment with homogeneous radios, heterogeneous radios, or both, and whether they are expected to react in real time.

Self-Organizing Cognitive Radios:

Self-organizing cognitive radios are the least sophisticated species to use spectrum awareness, and they may "think" only when they organize to form a network or modify it as more base stations of the same type are added or adjustments are made to improve performance. The behavioral template and rulebase is usually modified only by download of new software at intervals following initial network setup. Self-organizing radios usually assume that the other radios in the band are homogeneous, that is, operating using the same PHY and MAC framework. If other radios are present, they are assumed to behave consistently and have a high transmit duty cycle. Because of these simplifying assumptions, DRP is not a prerequisite and conventional radio technology with a "smarter" MAC approach can be used.

When presented with a spectral environment, these radios collect measurements on spectral activity and characteristics over a period of time, knowing what "signatures" to look for. The radio automatically institutes a discovery process to identify and contact other base stations that may participate in network formation, or isolate base stations that cannot participate thereby constituting mutual interference potential.

In this scenario, the self organizing radio carries a "Species memory". It knows that it is a "cellular" radio, as well as each cellular channel frequency and width. It must first identify existing large cell transmissions in each prospective channel to determine suitability for use in the underlay network without interference to the outside" system. It does this by "listening" for a period of time to determine if any transmissions are detected from embedded large-cell base stations, and assembles a histogram of signal activity over each channel in the band. After logging signal strength vs. channel over time, it can order channel reuse candidates by signal level. Using a programmed rule-base, any channel exhibiting peak signal level below a “safe-to-reuse” level is interpreted as “unused” in the area of the prospective nanocell, and made part of a candidate channel list. The candidate list is then ordered organized by ascending signal level.

The radio, with its smart MAC, then attempts to contact other base stations that may join the nanocellular network. It does this by sending messages via the backhaul network as well as over the air, using the list of candidate “white space” channels it has assembled. If other bases are contacted, a self-synchronizing protocol is employed to begin a series of reciprocal propagation loss measurements between the participating base stations. The process allows each participating base to transmit, wherein the others “listen”. When the last base has completed its transmission, the collected data is exchanged between the bases, and they cooperate to calculate prospective channel assignments and transmit power levels with the best reuse performance (lowest mutual interference). If insufficient “clear” channels exist to meet acceptable rule-base limits for adequate reuse separation, the bases again communicate to form the next-best formulation with two or more bases dividing the time resource between them. With channels and power levels chosen, the base stations go “on the air”. As clients make calls, signal strength of each one is monitored, and over a period of time the serving base station adjusts power levels to provide acceptable quality at cell edges.

Several research Labs have innovated techniques for self-organizing indoor cellular applications using the SDR base stations previously discussed as well as wireless LAN systems. Prototypes of these systems have indicated that self-organizing cognitive radio can be a powerful tool for reducing installation costs while providing stable network operation with good efficiency.

Adaptive Cognitive Radios:

Adaptive cognitive Radios take less for granted than self-organizing radios. Like Self-organizing radios they know their mission: Set up operation in a specified band either alone or with a group of other radios. But they do not assume that all radios may be of the same “species” or that the channel parameters are necessarily known. This is an important distinction, since it means that the "signatures" of other systems must be detected and identified before performing the cognition described in the last section. DRP's ability to tune, channel-filter, and demodulate arbitrary waveforms is ideal for such applications, in fact almost a prerequisite.

In an adaptive scenario, a radio may be seeking to operate in a band already used for other licensed or unlicensed services. As an example, let us assume that a municipal WiFi system seeks to operate in a band already populated by another service, perhaps licensed. The air interfaces and channel bandwidths may be very different. In such a situation, the radio may have to sample the entire band at various channel bandwidths to try to determine what transmissions are present. It must then attempt to identify whether it can operate in the same or a nearby channel without causing interference. After determining what channels might be used, it can act as a self-organizing cognitive radio as in the previous section.

Adaptive cognitive radios may also evolve to use the spatial reuse dimension. As advanced antenna processing techniques mature, these may be used to change the shape of the antenna pattern and power level to cooperate with existing services.

For example, the "reach" of an adaptive base station can be automatically adjusted to extend further in some directions compared to others to minimize interference to other services.

Adaptive radios must operate in an environment where other adaptive and self-organizing radios may already be operating. As a result, they may be required to modify their behavior as a result of more continuous monitoring than would normally be expected of self-organizing systems, which behave similarly to fixed operations except when the network re-configures due to addition of new base stations. Such monitoring may, for example, detect a new interferer that forces automatic system reconfiguration.

Dynamic Cognitive Radios:

Dynamic Cognitive Radios are breed-apart , and could be considered the top of the food chain" based on intelligence. They can not only perform as self-organizing and adaptive cognitive radios, they can act in real time to "scavenge" pieces of spectrum as a function of time. These radios may change their frequency, bandwidth, transmission timing, antenna pattern, and other parameters "on the fly" according to their perception of "white space" opportunities at a particular moment.

Thus, as other transmissions arise and disappear, the dynamic cognitive radio works with its MAC to "insinuate" its transmission into the environment produced by all others. Clearly maintaining continuous links between radios requires that the transmitters and receivers execute synchronized "handoffs" orchestrated by notification of the parameters for the next burst as the last one ends. Such complications beg standards activity to define new protocols to allow dynamic cognitive radios to interoperate.

The Current Cognitive Radio:

In 2002, the NIST Spectrum Policy Task Force Report[18] spotlighted cognitive radio with SDR as good candidate for enhancing spectrum utilization, while recommending that the FCC consider modifications of spectrum policy to leverage such capabilities. The report also contained studies that disclosed "white spaces" where spectrum appeared to be underutilized in some geographic areas.

In December of 2002, the FCC adopted an NOI in Docket 02-380 exploring unlicensed sharing of spectrum in TV bands below 900 MHz, observing that interference could be avoided by geographical separation with license database checking, or listen-before-talk algorithms. The waters were clearly being tested for possible cognitive radio techniques.

The NIST report and preparation of the NOI spurred heightened interest in cognitive radio at the FCC, and prompted an OET Workshop on Cognitive Radio Technologies in 2003. A series of presentations at the meeting spanned the range from why cognitive radio should be considered for spectrum management, through an analysis of spectrum "holes" and how to detect them, to full SDR and cognitive dynamic SDR air interfaces. One of the prominent contributors was DARPA, whose neXt Generation (XG) Communication program aims toward use of adaptive techniques to manage co-use of spectrum with multiple applications. The goals of XG are to boost spectrum utilization tenfold, and to facilitate global regulatory alignments and compliance. DARPA, with its Department of Defense affiliations, wishes to ensure that future military radio communication systems can work in concert with diverse international spectrum regulations and agencies, a significant and growing problem. As such, the program is not aimed so much at new radios themselves (although it assumes availability of advanced SDR platforms), but rather a group of advanced technologies for adaptive and dynamic cognitive radios.

The group envisions its work as providing an abstraction layer for managing flexible radios that could be used with many system architectures, developing a software protocol framework, and identifying core interference-related modules to which other XG capabilities can be grafted. To implement its program, it is crafting a "meta-language" to describe radio controls, conduct measurements, and command policy-based radio behaviors for particular locations and situations. Spectrum goals for XG are to use arbitrary pieces of new spectrum, use "bundled" fragments of spectrum, and "squeeze into existing used spectrum while not causing unacceptable interference.

Later in 2003, the FCC issued a Notice of Proposed Rulemaking and Order for "Facilitating Opportunities for Flexible, Efficient, and Reliable Spectrum Use Employing Cognitive Radio Technologies" and "Authorization and Use of Software Defined Radios". The NPRM discussed a wide range of cognitive radio issues, methodologies to encourage the technology, and proposed rule changes to hasten adoption. The FCC noted that the World Radio Congress meeting in 2007 has adopted an agenda item regarding use of cognitive radio as an extension to IMT-2000 systems, based on ITU-R activities.

For rural and unlicensed applications, a particular proposed rule change was interesting: Use of cognitive radio to increase maximum power levels in ISM and other bands in geographic areas appearing to have "limited spectrum use". These provisions would be most useful in wireless LAN and WISP services, as they could allow significant increases in "reach" (as much as 2.5 times). Such a change would provide an incentive to capitalize on cognitive radio technology in much the same way as higher power had been used to encourage use of spread-spectrum techniques in ISM bands earlier. The initiative is a good example of the shift toward "application-based" rulemaking.

A key barrier to extracting the benefit of these new rules would be, as it was with spread-spectrum interference, to determine how higher power levels would affect devices already operating in the band, how to detect the presence of these devices, and how to define "limited use" based on measurements. The sensing of the potential presence of other devices is one of the most difficult of envisioned cognitive radio functions, since it can mandate almost continuous SDR scanning, interpretation, and discontinuance of transmissions of the sensing radio if activity is detected over a predefined threshold.

The FCC thus began formulation of concepts for quantifying "activity" in a channel based on noise measurements over a period of time, and this work resulted in the "noise temperature" approach. This measurement was provisionally defined as measured aggregate noise plus interference power referenced to thermal noise. The temperature at which the channel was assumed "active" was defined as 30 dB above the calculated thermal noise floor in a measurement bandwidth of 1.25 MHz. A device would have to apply this measurement across a whole band to qualify for higher power benefits.

Even considering the seemingly straightforward noise temperature policy, many other questions still remained: Some devices that could be impacted might not transmit continuously (thereby complicating detection), or might be "held off" by legacy CSMA-like protocols with their own transmit threshold determinations. Moreover, characteristics of the survey antenna can play a strong role in the measurements, leading to possible misjudgments by cognitive radios. In VHF and UHF bands, changing propagation conditions could be problematic, as with day-night and seasonal variations.

Interestingly, the FCC NPRM also mentioned that it had issued a Notice of Inquiry soliciting comment on possible unlicensed use of TV broadcast spectrum in the 3650-3700 MHz band, indicating that such uses would probably require cognitive features to avoid interference. However, in a later Notice of Proposed Rulemaking (NPRM), "Unlicensed Operation in the TV Broadcast Bands," ET Docket No. 04-186, the FCC requested comments on contention-based protocols for use with cognitive radios, but did not expressly propose rules for such devices.

This restraint was probably indicative of its consideration of the barriers to harnessing the potential of cognitive radio: (1) Difficulty in diverging from traditional "centralized" regulatory processes, (2) slow emergence of cognitive radio algorithms that could be standardized and published, (3) proof that cognitive radio can work well in a broad range of situations. and (4) potential cognitive radio embarrassments, leading to devaluation of spectrum assets.

To provide a more circumscribed application space where cognitive radio could be put to use earlier with less liability, the FCC opened the possibility of so-called "secondary markets". These are markets where owners of systems could negotiate with others holding licenses for use or co-use of the asset. Spectrum leasing using cognitive radios might allow long-term or short-term leasing with end-of-lease relinquishment on an ad-hoc or real-time basis, for example by an exchange of tokens. Such arrangements could dramatically increase spectrum utility for more users.

A particular situation described in the NPRM was public safety spectrum, which is lightly-used normally, but heavily used in extraordinary conditions. Holders of this spectrum could "rent" it to other services, but recover it quickly in emergencies using reversion mechanisms such as "beacons". Beacons could be transmitted by the owner continuously as long as the spectrum could be used by the lessee. If the beacons ceased, the lessee would have to discontinue transmission until the beacons reappeared. The FCC pointed out that the beacons could be sent over wired as well as wireless channels.

The FCC also observed that coordinated spectrum sharing using cognitive radios might be possible between disparate licensed operations, for example satellite and terrestrial services.

A wide variety of interested parties commented on the NPRM, but few provided details on how such systems could be immediately implemented. One theme was repeated frequently among primary license holders: "we are concerned about uncontrolled interference in our spectral asset, and will require strong capabilities for remote interrogation and shutdown of suspected offending cognitive radios".

Reacting to FCC interest in cognitive radio, the IEEE began to investigate standards for a near-term opportunity that appeared to be a good proving-ground for cognitive radio: the clearly underutilized US television spectrum between 54MHz and 862MHz, which is being vacated (reluctantly) as broadcasters move to HDTV. The FCC has proposed release of up to 300 MHz of this UHF/VHF spectrum as its first major test of software defined or cognitive radios. It would permit fixed access systems transmitting up to 1W in power and portable devices up to 100mW. If successful, the

FCC could allow cognitive radios in other licensed bands to coexist as unlicensed devices, and encourage other regulators around the world to follow suit.

In late 2004, a new IEEE working group, 802.22, was chartered to develop a standard for a "cognitive radio-based PHY/MAC air interface for use by license-exempt devices on a non-interfering basis in spectrum that is allocated to the TV Broadcast Service". Another IEEE standards group, 802.16, argued that cognitive radio work should be under its aegis rather than in a separate group, but was defeated.

The first focus of the 802.22 effort is on Wireless Rural Access Networks (WRANs). The chair of the new group indicated that the spectrum is ideal for regional networks to provide broadband service in sparsely populated areas. 802.22's goal is to equal or exceed the quality of DSL or cable modem services, and to be able to provide service in areas where wire line service is not economically feasible due to the distance between potential users. In fixed networks, 802.22-based technologies could achieve multi-kilometer range and complement local Wi-Fi and 802.16 backhaul. 802.22 has succeeded in getting active and cooperative participation by the incumbent licensees (including their association MSTV), TV set manufacturers (through CEA), the IEEE Broadcast Technology Society, and other interested parties, as well as the "traditional IEEE 802 networking equipment and semiconductor producer participants.

The group indicated that its intent is not to "compete" with other IEEE 802 wireless networking solutions, but rather to develop a specialized solution to the particular application of making effective, non-interfering use of unused TV spectrum. However, it is clear that a successful implementation would probably propel use in other bands.

Other standards activities have also been opened to study processes for cognitive radio: The 802.16h License-Exempt Task Group, DySPAN (the 2005 IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks), Software Defined Radio Forum (SDRF), SIG1 Spectrum Topics (including WG4 New Air Interfaces, Relay-Based Systems and Smart Antennas), ITU-R WP8A Non Cellular and WP8F Cellular groups, and E2R, an Integrated Project of the 6th Framework Program of the European Commission, addressing the core of the strategic objective "Mobile and wireless systems beyond 3G".

Very recently, a new study group has also been opened in 802.11 for standardization of "Contention-Based Protocols for Cognitive-Radio Applications" in the 3650-3700

MHz band and bands below 900 MHz described in Report and Order and Memorandum Opinion and Order 05-56.

Clearly academia, standards, and industry have responded to the FCC's interest in using cognitive techniques to better-manage spectrum. But can the compelling nature of the opportunity be balanced against the practicalities of implementing such intelligent systems, and what are the steps from now to the future?

The Future of Cognitive Radio:

Like most applications involving “cognitive” computing, attempts to emulate reasoned thought usually prove to be much harder than at first envisioned. As we have discussed, earlier smart radios have already implemented elements of the stimulus-response state machine through the configuration and operation phases. These elements determine presence of a particular system, recommend the best operating channel to a base station coordinator, adjust power levels, and “gearshift”. However, bigger challenges surface as one moves toward true cognition: continuous monitoring, pattern matching (interpretation), and adaptation that close multiple feedback loops. It is the rich feedback that provides the capability for “learning” with process self-modification. Although learning provides much of the glamour of cognitive approaches, self-modifying systems are complex by nature.

The group indicated that its intent is not to “compete” with other IEEE 802 wireless networking solutions, but rather to develop a specialized solution to the particular application of making effective, non-interfering use of unused TV spectrum. However, it is clear that a successful implementation would probably propel use in other bands.

Other standards activities have also been opened to study processes for cognitive radio: The 802.16h License-Exempt Task Group, DySPAN (the 2005 IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks), Software Defined Radio Forum (SDRF), SIG1 Spectrum Topics (including WG4 New Air Interfaces, Relay-Based Systems and Smart Antennas), ITU-R WP8A Non Cellular and WP8F Cellular groups, and E2R, an Integrated Project of the 6th Framework Program of the European Commission, addressing the core of the strategic objective "Mobile and wireless systems beyond 3G".

Very recently, a new study group has also been opened in 802.11 for standardization of “Contention-Based Protocols for Cognitive-Radio Applications” in the 3650-3700 MHz band and bands below 900 MHz described in Report and Order and Memorandum Opinion and Order 05-56.

Clearly academia, standards, and industry have responded to the FCC’s interest in using cognitive techniques to better-manage spectrum. But can the compelling nature of the opportunity be balanced against the practicalities of implementing such intelligent systems, and what are the steps from now to the future?

The Future of Cognitive Radio:

Like most applications involving "cognitive" computing, attempts to emulate reasoned thought usually prove to be much harder than at first envisioned. As we have discussed, earlier smart radios have already implemented elements of the stimulus-response state machine through the configuration and operation phases. These elements determine presence of a particular system, recommend the best operating channel to a base station coordinator, adjust power levels, and "gearshift". However, bigger challenges surface as one moves toward true cognition: continuous monitoring, pattern matching (interpretation), and adaptation that close multiple feedback loops. It is the rich feedback that provides the capability for "learning" with process self-modification.

Although learning provides much of the glamour of cognitive approaches, self-modifying systems are complex by nature. Like any complex control system, a self modifying one may exhibit undesirable or destructive behaviors that do not always manifest until a particular set of conditions is presented. Stability can be more difficult to ensure as the number and diversity of possible inputs increases and the control system is asked to respond faster. If one considers cognition a continuum ordered by amount of input information and reaction time, one may consider self-organizing, adaptive, and dynamic cognitive radios as steps on an evolutionary ladder leading to increased "consciousness" and more complex environmental interactions.

As if the difficulty of designing "baseline" cognition is not hard enough, evolution can lead to additional instabilities if radios with differing levels of "consciousness" are mixed in the same environment. For example, populating a band with different kinds of "learning" radios operating with different perceptions of the environment and with differing "decision latencies" can cause instabilities not only between individual radios, but also between entire systems. Also challenging is institution of spectral "fairness" between competing cognitive systems using different algorithms, since a "greedy" algorithm may "freeze out" newcomers. Thus, the major issue with cognition is balancing the degree of "consciousness" with the liability of aberrant behavior and the probability that the behavior may occur.

Opening The Door to Cognitive Radio

Dr. Hossein Eslambolchi

March 1, 2012

In more recent proposed rulemaking and industry discussions, the FCC already appears to have tempered its initial enthusiasm with realization of the difficulties. It is choosing a few "test cases" in bands populated with fixed, non-cognitive, licensed radios operating with a good deal of legacy-rule "buffer" protection. In these situations, cognitive rules-based can be simplified, radio decisions can be "checked and balanced" by license database supplementation, and 2^{nd} order instabilities between radio types can eliminated, at least in the near-term, through tight specification of cognition parameters.

Using the added spectrum provided by the early cognitive radio test cases, new services housed in these bands can be combined with existing services and spectrum by using the multi-service radios described previously.

Today, the technologies best suited for the "test" bands are mainstream commercial (*e.g.* non-military) radio systems that have reached the "self-organizing" level of consciousness, where they can recognize their own kind and optimize system behavior in an environment also containing other non-cognitive radios. Even this small beginning is very valuable, as fledgling self-organizing systems are now at the heart of next-generation enterprise, municipal, and rural broadband network architectures using packet radios.

Examples of standards-based air interfaces with relevant fields-of-use would be 802.11 fortified with the newest QoS, Radio Resource Measurement, Wireless

Network Management, and Interworking with External Systems standards enhancements, 802.22, and 802.16 with a "cognitive" annex. These systems already

rely or will rely on OFDM transmission, which supports Fast-Fourier Transform (FFT) spectral analysis at least within the native channel bandwidths and modulation definitions covered by the standards. Some extra capability might have to be provided to be consistent with FCC noise temperature and occupancy definitions for operation in bands and channels other than those currently standardized.

The ability to use extensions of current hardware for the test cases is an important consideration, because it materially shortens time to availability and lowers cost. In modern "reciprocal" transmission systems that can use the same baseband radio processing in base station and client, it is usually client volumes that dominate radio cost. Thus, if existing clients can be economically modified to reach into new frequencies and cognition, systems are likely to flourish. This suggests that full DRP should not be a prerequisite for new rules in near-term test bands.

Likewise, If the spectrum released by the FCC for cognitive use is close enough to existing operating frequencies of the standards, complete redesign of RF front-ends may not be necessary, and it may be possible to cover existing and "cog" bands by selecting front-end filters depending on which band is in use. Using such technology "stretching" 802.11a could, for example, be made to operate in 6 GHz microwave spectrum, currently populated by legacy point-to-point microwave relay licensees.

Such a concept illustrates a particularly useful property of new small-cell outdoor networks leading to next-generation wireless access derived from today's wireless LANs. The size of the cells in these systems is small enough that underlays can provide a good balance of usefulness and interference avoidance in existing licensed and unlicensed bands exploiting possible higher transmit power via cognitive radio rules.

However, unlicensed bands may be occupied by many radio types and services. Areas with competing uses might not be well controlled (*e.g.* areas other than where municipalities could limit use of outdoor networks to a single choice). Larger cells also exacerbate the problem because interference liability can reach further and impact more users. It may be necessary to consider adaptive cognitive radio implementations instead for these cases.

Expanding Cognitive Radio Use:

The true future of adaptive cognitive radios is in bands with a less predictable mix of embedded radios. It carries with it two challenges: Development of behavior and adaptation rules that can cooperate with the wider variety of non-cognitive and self-organizing radios and availability of economical mainstream DRP processing platforms that can be frequency-agnostic and have enough SFDR.

Spurious Free Dynamic Range, the ratio of the maximum and minimum signals that can be received without distortion, is a major limitation of the DRP architecture. It is important in DRP receivers because the front-end and analog-to-digital converter must have sufficient SFDR to accommodate the peak-to-peak amplitude of the waveform represented by *all* of the signals in the band being processed. Accordingly, radios designed to remain linear with many (additive) strong signals, even far distant in frequency but in-band, can be limited in their ability to receive weak signals. For systems involving mobile clients, where wide dynamic signal strength swings are commonplace, radios must be designed with large SFDRs.

Currently DRP implementations are just entering the marketplace for commercial cellular base stations and military systems, but clients have not yet appeared except in prototype form. As DRP with high SFDR usually demands more battery power, those few clients that have appeared show some difficulty in reaching the battery life expectations users are accustomed to with cellular telephones and other portable appliances. This normally acts as a deterrent to viability of mainstream solutions. Moreover, as one moves from commercial to military uses, the portfolio of bands such radios might be expected to use also widens, mandating more complex front end designs that approach continuous tuneability over wide frequency ranges, further increasing bulk and cost.

One might expect that adaptive cognitive systems will first appear in bands where both cognitive radio and DRP are jointly specified for entirely new systems, such as next-generation satellite/terrestrial communication systems, where their precision, frequency agility, and ability to rapidly compensate for changing environments makes them more attractive for extremely wide-area mobile and portable use. Such applications might have to form overlays of a variety of other services, as in multi-band aggregations containing existing cellular, point-to-point, satellite uplink/downlink, and direct-to-satellite systems. Adaptive DRP-based systems are also better able to cope with the wide range of radio administration rules that continental or worldwide-scale systems would encounter.

Point-to-point (PTP) and Point-to-Multipoint (PTMP) 802.16d-like and 802.22 systems would also constitute a useful vehicle for adaptive cognition, since "clients" in these systems are actually fixed terminations. Applications with fixed link endpoints also makes adaptation less challenging. As such systems are usually contemplated for delivery of bulk broadband services in line-of-sight (LOS) or near-LOS situations, SFDR poses less of a problem because antenna gain and directivity can be used to limit the range of signals a DRP receiver might have to accommodate. Such systems are usually less cost-sensitive as they emphasize availability rather than single-user portability, and can better tolerate the cost of a more-expensive radio. Adaptive cognitive radios for this application could also be designed with a modular RF front-end architecture so that the band of operation could be selected upon installation rather than mandating a more-expensive multi-band "one-size-fits-all" approach.

Pushing The Envelope of Cognitive Radio:

When less-expensive processing and high-SFDR DRP solutions become available in low power, highly-integrated chipset form, it is natural to move into dynamic cognition. Dynamic systems represent the leading edge of radio agility and intelligence and, conversely, the biggest interference threat to existing services if they make mistakes or are delayed in relinquishing "white space" they have already "captured". Because they represent the fastest-response of all cognitive radios, and may have to operate with less-responsive generations of cog-radios in addition to non-cognitive radios, opportunities for instabilities are much more plentiful.

As a consequence, commercial use of dynamic cog-radios for the foreseeable future is probably relegated to licensed bands with well-known non-cognitive radios of particular design already operating, such as in cellular bands where highly periodic waveforms allow simple detection and synchronization. In such cases, one can envision dynamic radios that, like the early CDPD systems, "sandwich" their transmissions into open frequency, time, code, or spatial segments. A particularly fertile field for exploration is use of spectrum "holes" that are artifacts of frequencymultiplexed systems whose channel assignments were dictated by older hardware IF filters. Using DRP technologies, extremely discriminating filters can be realized that could allow transmission with minimal disturbance of existing services. Advanced OFDM techniques such as "water-filling" could be used in conjunction with multi-user interference suppression to adjust power of and "load" tones placed in the "holes" according to a template that is the inverse of the typical response of the older IF filters.

Out of the commercial sphere, dynamic cognition will remain the more likely domain of government and military radio communities who in the past were granted allocations but have long been less subject to specific FCC regulations. Recently though, the Spectrum Policy Task Force report has awakened the FCC to the possibility that much of the spectrum granted governmental and military communities may be very underutilized. Now that personal communication has intensified the appetite for spectrum (particularly some "beachfront" spectrum segments owned by these agencies), the FCC seems to considering "re-farming" these previously "untouchable" areas. The FCC's lead in rethinking such allocation policies in light of new technologies is being emulated in many World Radio Congress-represented countries. With governmental and military users being pressured to use more commercial communication hardware platforms instead of expensive special purpose designs, the combination of dynamic cognition and DRP is an attractive option if spectrum sharing with commercial systems comes to pass.

In terms of mainstream spectrum policy, as compelling as dynamic operation is, it will probably be sanctioned (officially) only in new "re-farmed" bands or segments of spectrum where the mix of legacy systems creates periodic, non-cognitive system signatures in frequency, time or both, rather than in situations where bursty, evanescent traffic might be encountered. In such cases dynamic radios can synchronize with each other or with existing transmissions, making use of the real-time capability and minimizing interference by remaining mostly-orthogonal in frequency, time, and space.

With clandestine or multi-asset war-fighter applications less encumbered by regulation, and where ability to operate in arbitrary spectrum, with virtually any combination of embedded systems, and with less concern regarding disruption of existing services, dynamic cognitive radio could become the first "general coverage" radio system capable of interacting with, or hiding from, any other radio. As such, it could be the "universal radio" engineers have contemplated since early radio regulation and design launched species differentiation.

Summary:

This paper has examined the history, concepts, types, and opportunities of cognitive radio technology, while forecasting how it might evolve against the backdrop of spectrum and regulatory considerations.

Cognitive radio is a compelling concept, and its possibilities capture the imagination of radio technologists as easily as the prospect of "thinking" robots[20] captures the imagination of all of us. The availability of inexpensive "brains", actuators, and sensors have direct analogs in modern radio "MACs", transmitters, and receivers, and both technology paths conjure visions of a completely redefined future where today's limitations are erased.

However, like the road to "human" robots, cognitive radio is a journey, not a destination. As compelling as the concept is, it likely to evolve slowly, and its impact on spectrum policy will undoubtedly be constrained by its ability to prove it can be

practical: It must show that it can extract substantial utility in shared spectrum without materially jeopardizing the value of current services we depend upon.

This is not to say that we should not embark upon the journey, however, because if only some of the potential benefits can be realized, the additional spectrum resource available to feed the growing demand for personal communications would be significant. Growing from the foundation of today’s technology this appears possible, and continuing advances in cognitive systems will only improve the potential.

Dr. Hossein Eslambolchi

March 1, 2012