



Cognitive Radio for Flexible Mobile Multimedia Communications

JOSEPH MITOLA III

The MITRE Corporation, McLean, VA 22102, USA and the Royal Institute of Technology (KTH), Stockholm, Sweden

Abstract. Wireless multimedia applications require significant bandwidth, some of which will be provided by third-generation (3G) services. Even with substantial investment in 3G infrastructure, the radio spectrum allocated to 3G will be limited. Cognitive radio offers a mechanism for the flexible pooling of radio spectrum using a new class of protocols called formal radio etiquettes. This approach could expand the bandwidth available for conventional uses (e.g., police, fire and rescue) and extend the spatial coverage of 3G in a novel way. Cognitive radio is a particular extension of software radio that employs model-based reasoning about users, multimedia content, and communications context. This paper characterizes the potential contributions of cognitive radio to spectrum pooling and outlines an initial framework for formal radio-etiquette protocols.

Keywords: software radio, cognitive radio, spectrum management, software agents

1. Background

1.1. Software radio and SDR

A software radio [1] is a multi-band radio capable of supporting multiple air interfaces and protocols through the use of wideband antennas, RF conversion. Analog to Digital Converters (ADCs) and DACs. In an ideal software radio, all the aspects of the radio (including the physical air interface) are defined in software on general-purpose processors. For some air interfaces such as Wideband Code Division Multiple Access (W-CDMA), such an ideal implementation may not be practical for one reason or another (e.g., power consumption). As processor technology advances, however, air interfaces that require Application Specific Integrated Circuits (ASICs) today may be implemented on general-purpose processors. The software-defined radio (SDR) therefore compromises the software radio ideal in order to implement practical high-performance devices and infrastructure with current technology. SDRs are implemented using an appropriate mix of ASICs, Field Programmable Gate Arrays (FPGAs), Digital Signal Processors (DSPs) and general-purpose microprocessors. The architecture and computational aspects of the ideal software radio have been defined formally [2]. In addition, the global SDR forum has defined an architecture framework, object models and other recommendations for SDR [3].

1.2. Cognitive radio

Cognitive radio signifies a radio that employs model-based reasoning to achieve a specified level of competence in radio-related domains [4]. Cognitive radio architectures being investigated at KTH employ the cognition cycle illustrated in figure 1.

The outside world provides stimuli. Cognitive radio parses these stimuli to recognize the context of its communications tasks. Incoming and outgoing multimedia content is parsed for the contextual cues necessary to infer the communications context (e.g., urgency). Thus, for example, the radio may infer that it is going for a taxi ride (with some probability) if the user ordered a taxi by voice and is located in a foreign country. The Orient-stage decides on the urgency of the communications in part from these cues in order to reduce the burden on the user. Normally, the Plan-stage generates and evaluates alternatives, including expressing plans to peers and/or the network to obtain advice. The Decide-stage allocates computational and radio resources to subordinate (conventional radio) software. The Act-stage initiates tasks with specified resources for specified amounts of time.

If the main battery has just been removed, however, the Orient-stage would *immediately* invoke the Act-stage to save data necessary for a graceful start-up after shutdown. Unexpected loss of carrier (e.g., because of departing a non-intelligent building) of an RF LAN that is in use can result in an *urgent* Decide-stage to restore traffic flow, such as via a more expensive 3G network. Most other *normal* events might not require such time-sensitive responses, resulting in the Plan-Decide-Act cycle of the figure. Cognitive radio also includes some forms of supervised and unsupervised machine learning.

Cognitive radio is a goal-driven framework in which the radio autonomously observes the radio environment, infers context, assesses alternatives, generates plans, supervises multimedia services, and learns from its mistakes. This observe-think-act cycle is radically different from today's handsets that either blast out on the frequency set by the user, or blindly take instructions from the network. Cognitive radio technology thus empowers radios to ob-

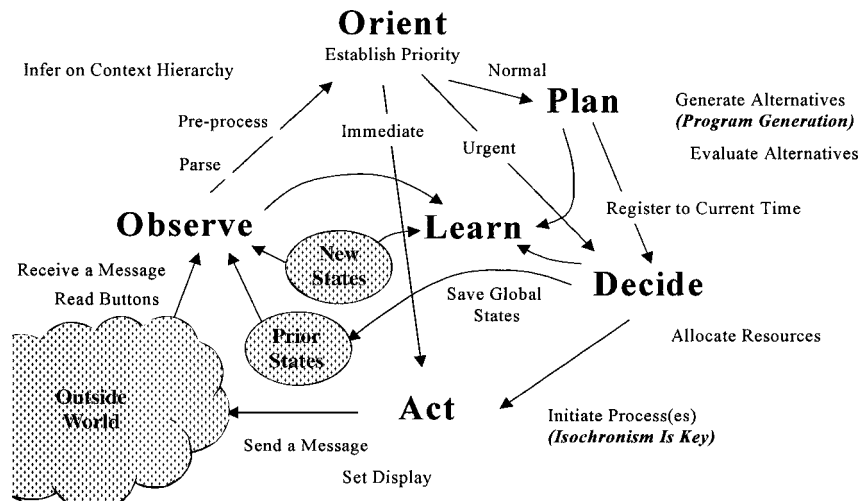


Figure 1. The cognition cycle.

serve more flexible radio etiquettes than was possible in the past.

1.3. SDK-enabled infrastructure and mobile multimedia devices

Recently Mitsubishi and AT&T announced the first “four-mode handset”. The T250 can operate in TDMA mode on 850 or 1900 MHz, in first generation Analog Mobile Phone System (AMPS) mode on 850 MHz, and in Cellular Digital Packet Data (CDPD) mode. This is just the beginning of the multiband, multimode, multimedia (M^3) wireless explosion. In the not-too-distant future, software-radio based Personal Digital Assistants (PDAs) could access a satellite mobile services, cordless telephone, RF LAN, GSM, and 3G W-CDMA. Such a device could affordably operate in octave bands from 0.4 to 0.96 GHz, (skip the air navigation and GPS band from 0.96 to 1.2 GHz), from 1.3 to 2.5 GHz, and from 2.5 to 5.9 GHz. Octave bands enhance antenna efficiency and reduce the cost of wideband RF conversion components. Not counting satellite mobile and radio navigation bands, such radios would have access to over 30 mobile sub-bands in 1463 MHz of potentially *sharable* outdoor mobile spectrum. The upper band provides another 1.07 GHz of *sharable*¹ indoor and RF LAN spectrum. This wideband² radio technology will be affordable first for infrastructure, next for mobile vehicular radios and later for handsets and PDAs. Such software radio technology expands opportunities for the dynamic sharing of spectrum. But it is the well-heeled conformance to the radio etiquettes afforded by cognitive radio that could make such sharing practical.

¹ Although $5.9 - 2.5 = 3.4$ GHz, only 1.07 GHz of this spectrum is allocated to mobile – hence *sharable* – subbands.

² Radios operating at 2.5 GHz are not expensive. Efficient wideband operation from 2.5 to 5.9 GHz, however, is expensive today. Costs are dropping as implementation technologies continue to advance.

2. Pooled radio spectrum

2.1. Pooling strategy

Present commercial wireless architectures are network-centric and constrained by spectrum allocations. Nick Negroponte [5] pointed the way with his spectrum management algorithm of the future: “If it moves, give it spectrum; if it doesn’t, give it fiber.” A slightly more practicable strategy accessible through SDR and cognitive radio would be to pool mobile spectrum:

“If it moves, give it spectrum pool precedence; if it doesn’t, make it pay.”

Satellites and aircraft move rapidly and/or cover large areas, so the bands dedicated to these vehicles would not be pooled. Broadcast television stations and 2 GHz microwave, however, would have to pay for the privilege of using prime spectrum in the middle of the mobile bands. The cost would be near zero in rural areas where the radio spectrum is uncrowded, so low cost broadcast and connectivity to rural areas would not be negatively affected. The cost would be appropriately larger in the densest urban areas where there is simply not enough suitable spectrum to economically meet the demand for mobile wireless. In addition, in those areas, the proliferation of fiber and cable makes such broadcast less necessary than in rural areas.

Conversely, however, those who currently “own” the spectrum could charge rent by the second, minute, or hour, as dictated by the marketplace. Federal, state, and local governments could generate revenue streams by literally renting channels that are not currently in use, and for which there is no need for some agreed-upon time interval (e.g., the next 10 minutes). The radio etiquette would specify that the cognitive radios would change bands and modes when legacy³ radios enter the bands. Thus, police forces would not have

³ A legacy radio is an existing non-cognitive radio operating in an air interface mode assigned to that band.

Table 1
Mobile spectrum pools.

Band	RF _{min} (MHz)	RF _{max} (MHz)	W _c	Remarks
Very low	26.9	399.9	315.21	Long range vehicular traffic
Low	404	960	533.5	Cellular
Mid	1390	2483	930	PCS
High	2483	5900	1068.5	Indoor and RF LANs

to procure new mobile radios. But, every police station, fire house, military facility, and taxi company could readily become an M³ cell site. These new base stations could pay for the SDR cell-site electronics with revenues from renting unused channels to 3G service providers during peak hours.⁴ Again, the radio etiquette can specify that emergency vehicles take precedence over commercial rental traffic. The cognitive radios using the bands would gracefully defer as they monitor the rented channels for FM radios, Tetra [6] users, and any other authorized legacy emissions. Since a variety of subbands and modulation schemes are in use by the public service community, the reliable identification of authorized users is sometimes not an easy task.

Jens Zander has pointed out that the shortage of spectrum can be viewed instead as a shortage of affordable cell sites [7]. Spectrum pooling both increases the number of sites and integrates the multi-mode spectrum so that the quantities needed for multimedia wireless applications are more affordable. The spectrum rentals (using e-cash or bartered spectrum channel-seconds, s-cash) would happen so fast that they must be accomplished by computer. Cognitive radios could use their knowledge of the RF environment, multimedia content, and communications context (e.g., a life-threatening emergency) to barter according to guidelines specified by the spectrum managers and represented in the radio etiquette protocol.

Clearly, managers set the criteria and make long-term spectrum leases. But the cognitive radios would rent out the short-term locally available spectrum that is not instantaneously in use, establishing spot prices as a function of time, bandwidth, interference levels, radiated power, location, and perhaps other parameters. Cognitive radios could rent spectrum for a second (e.g., to upload a brief email message), a minute (e.g., for part of a voice call), an hour (e.g., for a video teleconference) or more. Spectrum management authorities would establish the general etiquettes and constraints, but the market would set the price. Cognitive radios would use their spectrum awareness and goal-driven behavior to mutually assure conformance to etiquette, and to identify and report offenders to human authorities. Such a pooled spectrum strategy could accomplish dynamically with distributed control technology what could not be contemplated with today's centralized allocation-based control – the cost-effective efficient pooling of formerly scarce mobile radio spectrum for fair and equitable use when and where needed.

⁴ The resulting relative glut of spectrum could drive costs down in the short term, enhancing profit margins to provide the capital necessary to build out today's hardware-constrained infrastructure to M³-capable SDR infrastructure.

The resulting spectrum management process, however, would be akin to a chess game where the board is the radio spectrum, the players are the cognitive radios and networks, and the winners are the users. To see how the game might be played, the next section examines the “chess board” – the potentially pooled mobile radio spectrum. Subsequent sections describe the “legal moves” – specified by the radio etiquette protocol needed to maintain equity and order.

2.2. Pooled mobile spectrum parameters

The laws of physics impose limits on the spectrum that is useful for pooled terrestrial mobile-multimedia applications. The HF bands and below, for example, have limited bandwidth (practically speaking, a few tens of kHz, although heroic experiments have achieved megabit per second data rates in truly unique circumstances). HF also propagates for thousands of miles, a range that exacerbates cochannel interference. Bands above 6 GHz rely on directional antennas for reasonable data rates at reasonable ranges. HF and upper SHF are therefore not suitable.

The four spectrum pools of table 1, however, are ideally suited to mobile applications. The *very low band* of this mobile spectrum regime penetrates buildings and propagates well in rugged terrain. The *low band* has the best propagation for high-speed terrestrial mobile traffic, in part, because auto and rail traffic is supported with relatively low infrastructure density. The *mid band* is best for Personal Communications Services (PCS) with its higher infrastructure density. In addition, the high band has the large coherent bandwidth for high data rate Internet and mobile video teleconference applications. 3G waveforms could be used in any of these bands, but are best suited for the *low* and *mid bands*. W_c is the total spectrum that could participate in the spectrum pool based on an analysis of US, Canadian, and UK spectrum allocations [4]. W_c does not include satellite, aircraft, radio navigation, astronomy, or amateur bands, which are not suited for pooling. The pooling concept is illustrated in figure 2.

A cognitive radio-access network could evolve the operating parameters shown in table 2. For simplicity, the entire population offers load (100% penetration). With pooling, each mobile outdoor user would have an average of 432 kbps. This assumes today's infrastructure density and 2G-equivalent bandwidth efficiency (0.2 Mbps/MHz/cell). These rates are gross data rates not discounted for Quality of Service (QoS), which can lower these rates substantially if low bit error rates are required (e.g., for file transfers) [8]. They also do not include signaling overhead. On

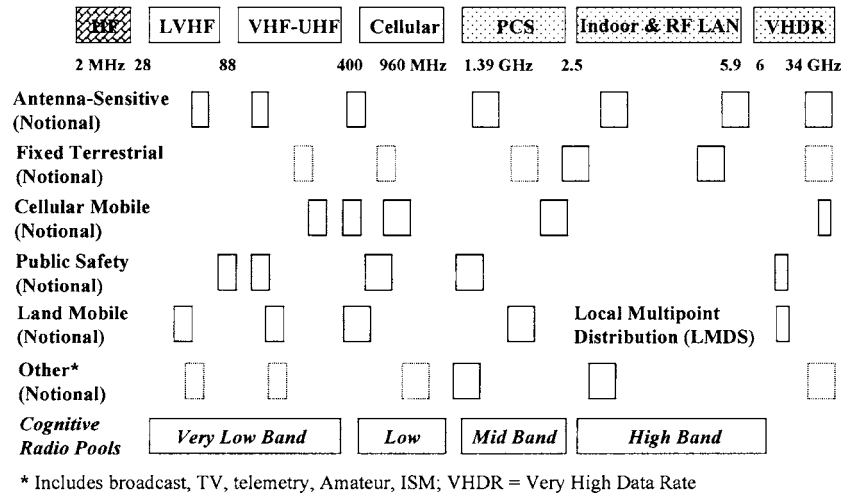


Figure 2. Fixed allocations versus pooling with cognitive radio etiquette.

Table 2
Illustrative cognitively pooled radio access network parameters.

Parameter	Illustrative range of values	Remarks
Total spectrum	0.4–2.5 GHz	1.463 GHz pooled
Duplexing	Frequency Domain (FDD)	Evolved from cellular services
Voice channel	8 1/3 kHz-equivalent, TDMA or CDMA	Evolved from second generation
Channels per cell	25088	Usable, including 6:1 reuse and FDD
Coverage area	4000 square kilometers	The size of Washington, DC
Number of cells	40 commercial (plus 40 public sites ^a)	5.5 km average cell radius (3.9 km)
Population	609,000	The entire population of Washington, DC
Offered demand	0.1 Erlang	Multimedia level (vs. 0.02 for voice user)
Demand per cell	1522 Erlang	Drops to 761 considering public cell sites
Spectrum per user	160.7 kHz (320.4 kHz with public sites)	0.22–0.64 Mbps/user (0.4–1.28 Mbps)

^a Public sites are towers of police, fire, military, and other government/ public facilities pooling spectrum.

Table 3
Multimedia implications of pooled spectrum.

Feature	Implications
Sharing of narrowband channels	Reduces costs of new uplink channels for Internet browsing and other multimedia services
Spectrum block rentals (e.g., 1–5 MHz)	New spectrum for multimedia downlinks
Increased infrastructure density	Greater availability of affordable multimedia
Short term spectrum rentals	Accommodate peak demand; Introduce new services incrementally

the other hand, 3G technology is supposed to achieve 0.45 Mbps/MHz/cell, so the rates are representative of the range of rates achievable with a mix of 2G and 3G technology.

With spectrum pooling, then, multimedia bandwidths can be achieved without a major increase in the number of cell sites. In part, the participation of public facilities increases the number of sites. In addition, the pooling of spectrum is more efficient than block allocations. Through cognitive radio etiquette, police, fire, and rescue units participating in spectrum pooling will have precedence for spectrum. They can also communicate seamlessly via the shared SDR cell sites.

2.3. Multimedia implications

Pooled spectrum is attractive for wireless multimedia applications for the reasons listed in table 3.

The benefits seem attractive. What about guarantees and fairness of use? In order for the approach to work, all emergency services, government functions, private, and commercial users will participate actively or passively a cognitive radio etiquette protocol that makes the difference between new levels of spectrum efficiency and chaos.

3. Cognitive radio etiquette

Radio etiquette is the set of RF bands, air interfaces, protocols, spatial and temporal patterns, and high level rules of interaction that moderate the use of the radio spectrum. Etiquette for spectrum pooling includes the spectrum renting process, assured backoff to authorized legacy radios, assured conformance to precedence criteria, an order-wire network, and related topics.

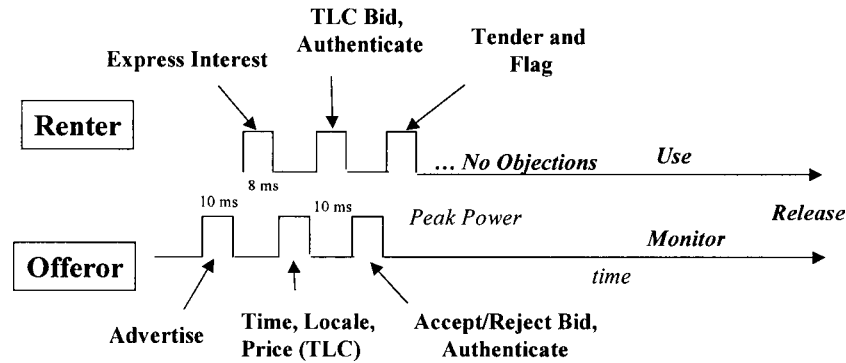


Figure 3. Time line of spectrum rental protocol.

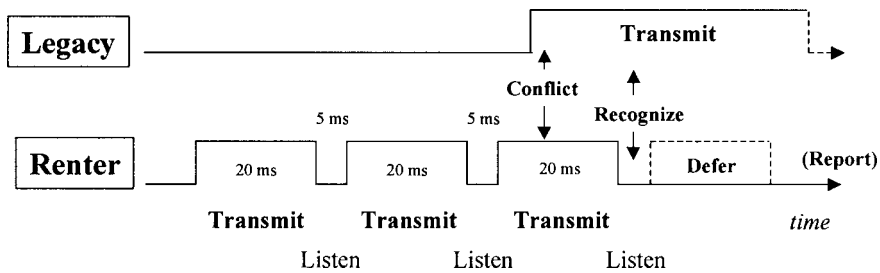


Figure 4. Polite backoff protocol assures channel access.

3.1. Renting spectrum

An initial protocol framework for renting radio spectrum is illustrated in figure 3. The time line shows the power levels in the rented channel, differentiating signals of renter and owner. The offeror initiates the process by posting an “Advertise” flag in the channel that is for rent. This in-band signaling accomplishes multiple goals. First, it unambiguously identifies the frequency, bandwidth (through its spectrum occupancy mask), and spatial extent of the channel (through the propagation of the signal). This signal should be pseudo-random, coded so that signal-processing gain can recover the signal when it is weaker than the noise and interference. A filtered PSK PN sequence of 100 bits duration at a 10 k chips per second rate, filtered to an 8 1/3 kHz bandwidth would advertise an 8 1/3 kHz channel in a 10 ms burst. The offeror listens for 10 ms and then repeats the advertise-signal twice more. The sequence starts as close as practicable to the tick of the offeror’s local GPS-second clock. The three-flag series repeats on the next second. A legacy user (or spectrum manager) could hear these bursts and realize that the channel is available for rent, expressing an objection by keying a transmitter.

A renter can express interest with a coded interest-burst similar to the advertise-burst. If the offeror hears the interest-burst, the second burst specifies the rental time interval, operating locale, and price of the channel. This data exchange would be Huffman coded using a-priori knowledge. The renter then submits a bid with a short authentication sequence. The offeror may accept the bid, authenticating itself in return. Finally, the renter tenders the (e- or s-) cash, completing the initial rental protocol.

Both then wait and listen for the rest of this GPS second for objections. The 100-bit objection-sequence would be nearly orthogonal to the advertise- and interest-sequences. A legacy radio that begins to use the channel in native mode (e.g., FM push to talk) automatically negates the rental agreement if its received signal strength at the renter or offeror location exceeds a threshold. A 100-bit rental-cancellation sequence from either party then cancels the deal. The offeror cannot attempt to rent the channel again until after a specified waiting period (e.g., seconds to minutes), or until re-advertised by the offeror. After using the channel successfully, the renter provides the additional e-cash validation bits required to secure payment of the offeror’s final bill. The initial rental protocol identifies the renter sufficient to pursue a claim if the renter defaults after using the channel. To avoid problems with the one-second granularity of the rental agreements, a service provider provisions the network by renting a few standby channels for traffic that cannot wait for the next rental period. During the use of the channel, both offeror and renter use a polite backoff protocol.

3.2. Polite backoff protocol: defer to authorized legacy users

If the cognitive radios using the channel employed a conventional air interface, legacy users would be unable to break in. Thus, for example, a police officer could not use his assigned frequency to call for assistance. The polite backoff protocol illustrated in figure 4 solves this problem, albeit at the expense of some loss of throughput.

The renter supplies (digital) traffic to the channel for 20 ms as shown in the time line. Both renter and offeror listen during the subsequent 5 ms listen-window. The high

rate of listen-windows assures that not more than 25 ms of legacy speech would be truncated, a level that should be essentially imperceptible distortion of a push-to-talk radio signal. If a legacy waveform exceeds a carrier to interference ratio (CIR) threshold, for either renter or offerer, the conflict is recognized and the channel is immediately vacated. The truncating party issues a 10 ms termination-burst that indicates the cause is legacy interruption. The rest of the traffic would be sent on another channel. Clearly, the renter can claim that the goods were not delivered, and not send the final payment bits. The renter and the offerer would log the time, place and other parameters of the legacy use of the channel so that spectrum managers could identify abusers. Since e-mail, attachments, file transfers, audio clips, video clips, and other asynchronous multimedia are relatively insensitive to end-to-end delay, this mode would be acceptable for low-cost wireless access. Of course, the offerer could make the channel available with peek-through required in only every N th listen-window, introducing some clipping at the onset of reclamation of the channel. This enhances the throughput (and price, perhaps) accordingly. Police might be able to offer only the high listen-window mode. Fire departments or military users might be willing to key the microphone for a second or more to reclaim the channel. Once reclaimed, the channel remains dedicated to the legacy user either until the expiration of a pre-defined time-out or until the channel is again offered for rent. In addition, the full protocol would contain sequences that indicate the channel has already been rented (a “sold” sign).

3.3. Precedence and priority

All users want guarantees that spectrum will be available when and where needed. Thus, any workable pooled spectrum approach has to have a way of providing such guarantees. Figure 5 provides an example precedence of spectrum uses. The character of the existing band allocations defines the default spectrum-use precedence. Designated authorities may change precedence globally or locally. The etiquette allows one to designate a user (e.g., by international mobile subscriber identification), a channel, or any combination of [user \times time \times space \times frequency] with a specific precedence.

If these notions are subjectively acceptable, the task remains to formalize them so that the radio control algorithms will perform as intended. In particular, the radios have to be able to infer many aspects of precedence from events. This is a technical challenge that requires that the radios become context-aware. The formalism therefore must provide some strong evidence that the system as a whole supports the statutory guarantees in spite of the sharing arrangements. One mechanism that supports such guarantees is an order-wire for the coordination of needs and plans for spectrum use and assignment.

1. **Emergencies** – Established by authorities, inferred from events.
2. **Government** – Attributed by band or channel modulation.
3. **Public Interest** – Default by band, inferred from events.
4. **Commerce** – Default by band and mode, inferred
5. **Other** – Recreational, sports, hobbies, etc.

Figure 5. Precedence of spectrum use.

```
{Tell :language RKRL :ontology Spectrum Rental
  :Rental_offer (:Owner Fairfax_Police :Location Chantilly_VA
    (:RF_low 451 :Nchannels 12 :From 141118 :Until 141523)
    :Allowed_formats ((:DECT 32kbps) (:GSM GPRS)
      (:Equivalent))
    :Legacy 25kHzFM)}
```

Figure 6. KQML expression of a plan.

3.4. An order-wire system and knowledge exchange language

The spectrum offerer may post an order-wire channel. An order wire is an ad hoc signaling and control channel. If the offerer does not post one, the cognitive radios using the band could create one using a peer network in which the first user becomes the network control station (e.g., JTIDS [9]). The details of such a network are not central to cognitive radio research, but the language used to represent general world knowledge, plans, and needs is a key issue.

The Radio Knowledge Representation Language (RKRL) is the language and knowledge structure being used to develop cognitive radio at KTH. The Knowledge Query and Manipulation Language (KQML) [10] was explicitly designed to facilitate the exchange of such internal knowledge. Based on “performatives” such as “tell” and “ask”, KQML readily express pooled-spectrum management information. KQML’s “:content” tag, for example, delivers unstructured content. Although this general purpose tag would suffice, the introduction of new tags for spectrum pooling imparts additional structure to the dialog. The new tags include :Rental_offer, :RF_low, :Nchannels, allowed_formats, :Legacy, :Equivalent, and tags for standard PCS formats such as DECT, GSM, and new 3G modes. The tags :From and :To refer to the time at which the rental is being offered. Using KQML, mobile nodes and networks may share plans about anticipated needs for spectrum so that it may be efficiently identified and rented. The KQML plan to offer spectrum uses the Tell performative to tell the (cognitive) network its plan as shown in figure 6.

The ontology performative could invoke a special format for the “:content” part of a normal KQML message. The tagged format shown here is meant to suggest both a more widely endorsed set of standard tags, and the opportunity for significant data compression on such messages. The network uses the ontology to look up defaults and compressed codes for the general knowledge expressed in the packet. The other aspects of the plan are self evident, which is one of the strengths of KQML.

3.5. Related topics

Cognitive radios need metrics for cost and value along with other high level rules of etiquette, such as:

- “tell the truth about who you are and what you need”,
- “block calls entering a disaster area and expedite those leaving”, and
- “test evolving protocols in off-peak hours”.

3.6. The complexity of a society of cognitive radios

Given a reasonable development of the spectrum pooling framework outlined above, one might think that performance could be projected using contemporary radio-engineering techniques. But cognitive radios will have complex internal structure and an ability to adapt to local circumstances. Since they also will be richly interconnected, they will form a complex adaptive system [11]. Cognitive radios could behave like an ant colony [12], evolving their own paths through the spectrum and intervening nodes to ferry voice, data, video, and multimedia packets through the ether. These radio-ants might move packets from the smaller power-starved radios through multimode vehicular radios and on to conventional cell sites. If the etiquette is too strict, very little additional benefit will come from spectrum pooling because the control overhead will be too high to be workable. If, on the other hand, the etiquette is too liberal, there will be much interference and universally poor quality of service. Such complex adaptive systems operate best “at the edge of chaos” [11]. This is not a particularly comfortable place for spectrum managers. Such complex adaptive systems are in fact difficult to understand, model and diagnose [13]. Nevertheless, they also produce efficient answers to NP-hard problems [12]. Thus, in some sense an ant-colony of cognitive radios left to evolve spectrum use among themselves might be the most efficient way to achieve high value from limited radio spectrum in a reasonable time. How can we structure the capabilities and etiquettes of cognitive radio so that spectrum pooling is workable? There are many important aspects to this question. Cognitive radio research at KTH continues to address such questions.

4. Conclusion

Software radios provide a vast untapped potential to personalize services. But the contemporary process of spectrum allocations takes years to decades and lacks flexibility. In part, this is because there is no reliable technology for guaranteeing spectrum use to its primary owners. This limits the flexibility and responsiveness of the radio to the network

and to the user. Cognitive radio offers the opportunity to employ spectrum rental protocols in a way that is sensitive to users and to the communications context. The cognitive radio rental etiquette is thus offered as an approach to more efficient use of a limited resource that is in high demand. Its agent knowledge and inference mechanisms are under development, as is the initial critical-mass definition of RKRL. The use of KQML in the exchange of plans among cognitive radios appears promising. The goal of the cognitive radio research is to develop software agents that have such a high level of competence in radio domains that they may accurately be called “cognitive”. The present research is merely a small step in this interesting research direction.

Acknowledgement

The author would like to thank Professor Gerald Q. Maguire and Professor Jens Zander of KTH, and Dr. Erland Wikborg of Ericsson for their insightful support of the research reported in this paper.

References

- [1] J. Mitola, *Software Radio: Wireless Architecture for the 21st Century*, 103255.2507@compuserve.com, Mitola's STATISfaction, ISBN 0-9671233-0-5 (1999).
- [2] J. Mitola, Software radio architecture: a mathematical perspective, *IEEE Journal on Selected Areas in Communications* (May 1999).
- [3] SDR Forum Web Site, www.sdrforum.org (1999).
- [4] J. Mitola, Cognitive radio – model-based competence for software radios, Licentiate Thesis, KTH, Stockholm (September 1999).
- [5] N. Negroponte, MIT Multimedia Laboratory, Cambridge, MA (1990).
- [6] Tetra is a radio technology of Ericsson, Stockholm, Sweden, www.ericsson.se.
- [7] J. Zander, Radio resource management in future wireless networks: requirements and limitations, *IEEE Communications Magazine* (August 1997).
- [8] J. Mikkonen, *Quality of Service in Radio Access Networks* (Tampere University of Technology, Tampere, Finland, May 1999).
- [9] R. Ziemer and D. Petersen, *Digital Communications and Spread Spectrum Systems* (Macmillan, New York, 1985).
- [10] T. Finin, KQML – a language and protocol for knowledge and information exchange, in: *International Conference on Building and Sharing of Very Large-Scale Knowledge Bases*, Tokyo (December 1993).
- [11] S. Kaufman, *At Home in the Universe* (Wiley, New York, 1995).
- [12] Dorigio and Gambardella, Ant colony systems: cooperative learning of the travelling salesman, *IEEE Transactions on Evolutionary Computation* (April 1997).
- [13] D. Fogel et al., Inductive reasoning and bounded rationality reconsidered, *IEEE Transactions on Evolutionary Computation* (July 1999).

Joseph Mitola III. Photograph and biography not available at time of publication.

E-mail: jmitola@mitre.org